

**USING INFORMATION TECHNOLOGY TO SUPPORT THE
DISCOVERY OF NOVEL KNOWLEDGE IN ORGANIZATIONS**

by

TRACY A. JENKIN

A thesis submitted to the School of Business

In conformity with the requirements for

the degree of Doctor of Philosophy

Queen's University

Kingston, Ontario, Canada

(August, 2008)

Copyright © Tracy A. Jenkin, 2008

ABSTRACT

In this dissertation, I examine how IT can support individuals, and in turn their organizations, in learning about and knowing their external environment on the Web. Specifically, I examine novel-knowledge discovery in the context of the multi-level organizational learning process, focusing on cognitive developments and changes to mental models. Novel knowledge is defined as knowledge that is potentially strategically important to the organization, not currently known to the organization, indirectly relevant and therefore difficult to find. Novel knowledge is proposed to be one of three different types of knowledge that organizations seek to discover in their environment. A theoretical framework is developed to identify the sets of tool characteristics, collectively referred to as levels, which are proposed to support the discovery of different types of knowledge, as well as different modes of learning and learning processes. In addition, extensions to the 4I organizational learning process model are proposed, specific to searching and learning on the Web: 1) adding a fifth process – information foraging and search-term development, and 2) adding a fourth level to the learning process – the machine-level. A competing theories approach is used to develop a rich understanding of knowledge discovery and learning on the Web. Understanding which types of tools are useful in different learning contexts has implications for learning effectiveness and may help firms understand how to “manage” their learning. Tools for the discovery of highly novel knowledge are less prevalent than tools to support the other levels of knowledge discovery. Accordingly, a design theory for novel-knowledge discovery tools is proposed based on organizational learning theories. An instantiation of the design theory, a novel-knowledge discovery tool, is developed and tested within the organizational learning process and compared to tools at the other two levels of knowledge discovery. In addition, different processes involved in using a novel-knowledge discovery tool at the group level are examined. Three separate studies were conducted, including

a lab and field experiment, and case study. The results are proposed to demonstrate how novel-knowledge discovery tools can support organizational learning.

ACKNOWLEDGEMENTS

After spending countless months researching, writing and defending a 469-page thesis, one would think that writing the acknowledgements section would be easy. However, expressing my sincere thanks to the many individuals who helped and supported me in this endeavor is a daunting task. In spite of the length of the list below, undoubtedly I have missed several individuals who helped in this process. To anyone I have left out, I apologize and give my thanks. Given my background in project management it is fitting that I compare my thesis to a major project with phases, deliverables, milestones, gates, constraints, issues, risks, costs, benefits, stakeholders, team members and resources. I was merely the project manager of this project and depended upon a large number of people to help me get the work done.

I would like to thank my committee members – Brent Gallupe and Jim McKeen – and especially my supervisors – Yolande Chan and David Skillicorn – for their help in shaping this dissertation and challenging my thinking. A huge thanks to Yolande and David for their direction, support, guidance and willingness to supervise a research project that was highly novel (and thus, complex) in so many ways. Without them, this research would not have been possible. I would also like to thank Bob Crawford – my internal-external examiner – and Betty Vandebosch – my external examiner – for their insightful feedback and willingness to serve on my committee.

Prior to starting my empirical studies, I attended the ICIS doctoral consortium in Milwaukee where I had the pleasure of working with faculty mentors Ann Majchrzak and Alan Dennis. I thank both Ann and Alan for their extremely helpful feedback.

I would like to sincerely thank Jim Provost for all of his hard work in developing Athens 2.0 across all three studies. Without Jim, this research would not have been possible.

Throughout this research, a number of individuals within Queen's School of Business helped me by providing resources and support – financial and technical. Specifically, I would like to thank The Monieson Centre, the Queen's Executive Decision Centre (Erik Lockhart and George Jackson), the QSB Information Technology department, Kathleen Lewis, the QSB Research Subject Pool, the QSB PhD office, and of course the students and organizations who participated in my studies.

A number of individuals helped me with several pretesting and review activities. Thanks to all of you. Specifically, I would like to acknowledge the support of Ana, Ann-Frances and Jim, who each went above and beyond.

Throughout my graduate studies, I have had the fortune of having a “loud” cheering section – people who have encouraged and supported me, even if they weren't quite sure what I was doing and thought I was a little bit crazy. Thanks to all of you – especially Greg Stott and Wolf Tormann.

I would also like to acknowledge and thank my family. To my parents – Keith and Carol Jenkin – thank you for all of your support, encouragement and riding the roller coaster with me. To my faithful and fearless fur-kids – Max and Sam – thank you for the millions of “bouche kisses”. And most of all, to my husband – soon-to-be-doctor Keith Rogers – thank you for everything you do.

TABLE OF CONTENTS

ABSTRACT.....	ii
ACKNOWLEDGEMENTS.....	iv
TABLE OF CONTENTS.....	vi
LIST OF FIGURES	xiv
LIST OF TABLES	xv
CHAPTER 1: INTRODUCTION	1
1.1 Novel Knowledge – A Case Example.....	1
1.2 Introduction to Novel-Knowledge Discovery.....	1
CHAPTER 2: LITERATURE REVIEW	6
2.1 Knowledge	6
2.2 Novelty.....	7
2.3 Novel Knowledge	10
2.4 Creativity	13
2.4.1 Person.....	13
2.4.2 Process	14
2.4.3 Product	15
2.4.4 Press (Environment).....	15
2.4.5 Groups and Creativity	15
2.4.6 Relationship between Novelty, Creativity and Innovation	16
2.5 Knowledge Discovery.....	17
2.6 Tools for Knowledge Discovery on the Web.....	22
2.6.1 Data Mining	22
2.6.1.1 Linguistic Analysis	24
2.6.1.2 Co-occurrence Analysis	24
2.6.1.3 Clustering and Categorization:.....	25
2.6.1.4 Visualization	25
2.6.2 Web Search Engines	25
2.7 Organizational Learning	27
2.7.1 Mental Models	28
2.7.2 Learning Modes	29

2.7.3 Individual vs. Organizational Learning	32
2.7.4 Organizational Learning Process	33
2.7.5 Learning Novel Concepts	38
2.7.6 Learning and IT.....	41
2.8 Summary: Integrating Knowledge, Learning, and IT	42
CHAPTER 3: THEORETICAL FRAMEWORK.....	43
3.1 Introduction.....	43
3.2 External Knowledge-Discovery Theoretical Framework	43
3.3 Competing Theory	55
3.4 Summary	57
CHAPTER 4: NOVEL-KNOWLEDGE DISCOVERY DESIGN THEORY	58
4.1 Introduction.....	58
4.2 Design Science and Theory	58
4.3 NKD Design Theory Components.....	62
4.3.1 Meta-Requirements and Meta-Design	63
4.3.1.1 Problem 1: Difficulty Locating Novel Knowledge.....	63
4.3.1.2 Design Principles for Problem 1	65
4.3.1.3 Problem 2: Difficulty Assessing Relevance of Novel Knowledge	66
4.3.1.4 Design Principles for Problem 2	67
4.3.1.5 Problem 3: Difficulty Learning Novel Knowledge.....	68
4.3.1.6 Design Principles for Problem 3	69
4.4 Evaluation of a Design Theory	70
4.5 Design-Theory Process	76
4.6 Summary	78
CHAPTER 5: GROUP PROCESSES AND OUTCOMES	79
5.1 Introduction.....	79
5.2 Overview: Group Processes, Support Systems and Outcomes	79
5.3 Outcomes	82
5.3.1 Radicalness	82
5.3.2 Benefit.....	83
5.4 Processes.....	84
5.4.1 Novel Knowledge and Group Challenges.....	84
5.4.2 Individual Use Only.....	86

5.4.3 Individual Use and GSS Tool	86
5.4.4 Individual Use and Collaborative-Filtering Tool.....	88
5.4.5 Comparison of Use Processes	90
5.5 Summary	92
CHAPTER 6: RESEARCH DESIGN.....	93
6.1 Introduction.....	93
6.2 Experiments	94
6.2.1 Overview.....	94
6.2.2 Key Design Elements of Experiments	96
6.2.3 Research Design.....	98
6.2.4 Lab Experiment.....	99
6.2.5 Field Experiment.....	100
6.3 Case Study	103
6.3.1 Overview.....	103
6.3.2 Research Design.....	104
6.4 Summary	106
CHAPTER 7: STUDY 1 – LAB EXPERIMENT.....	107
7.1 Purpose.....	107
7.2 Research Design.....	109
7.2.1 Measures	110
7.2.1.1 Independent Variable	110
7.2.1.2 Dependent Variables	110
7.2.1.3 Manipulation Check.....	111
7.2.1.4 Experimental Controls	112
7.2.1.5 Covariates	112
7.2.2 New Measurement Development.....	113
7.2.3 Card Sort	115
7.2.3.1 Phase 1 Card Sort.....	116
7.2.3.2 Phase 2 Card Sort.....	118
7.2.4 Survey Development.....	120
7.2.5 Experiment Pretest	120
7.2.6 Experiment Pilot Test	120
7.2.7 Experiment Execution.....	121

7.2.7.1 Sample.....	121
7.2.7.2 Pretest.....	121
7.2.7.3 Training.....	122
7.2.7.4 Task.....	122
7.2.7.5 Posttest	123
7.3 Analysis	124
7.3.1 Preliminary Analysis.....	124
7.3.1.1 Missing Value Analysis	125
7.3.1.2 Outlier Analysis	125
7.3.1.3 Normality Assessment	126
7.3.1.4 Exploratory Factor Analysis	127
7.3.2 Data Analysis	127
7.3.3 Idea Evaluation	127
7.4 Results.....	129
7.4.1 Exploratory Factor Analysis	129
7.4.1.1 Knowledge-Discovery Types EFA	129
7.4.1.2 Knowledge-Discovery Types: Scale Reliability	132
7.4.1.3 Learning Process EFA	133
7.4.1.4 Learning Process: Scale Reliability	134
7.4.1.5 Learning Mode EFA	135
7.4.1.6 Learning Mode: Scale Reliability	137
7.4.1.7 Full Factor Analysis.....	137
7.4.1.8 Summary: Scale Reliability	141
7.4.1.9 EFA Sample Size	142
7.4.1.10 PIIT, GCSE and WSE EFA	143
7.4.1.11 PIIT, GCSE and WSE: Scale Reliability	144
7.4.1.12 Summary	144
7.4.2 MANCOVA.....	145
7.4.2.1 Summary	154
7.4.3 Experimental Controls	155
7.4.3.1 Summary	156
7.4.4 Qualitative Results.....	157
7.4.4.1 Summary	164

7.4.5 Ideas	165
7.5 Discussion.....	165
7.6 Summary	170
CHAPTER 8: STUDY 2 – CASE STUDY	171
8.1 Purpose.....	171
8.2 Research Design.....	171
8.3 Case-Study Visit	172
8.4 Analysis	174
8.5 Results.....	176
8.5.1 Knowledge Challenges	176
8.5.2 Novel Knowledge	177
8.5.2.1 Importance	178
8.5.2.2 When and Why	178
8.5.2.3 Types.....	179
8.5.3 Novel-Knowledge Discovery.....	179
8.5.3.1 Processes	179
8.5.3.2 Types of Searching	182
8.5.3.3 Outcomes	183
8.5.3.4 Sources.....	184
8.5.4 Innovation Approaches	185
8.5.5 Athens Tool-Use	185
8.5.5.1 Positive Tool Experiences.....	185
8.5.5.2 Types of Problems	186
8.5.6 Learning Processes.....	189
8.5.7 Secondary Athens Evaluation	197
8.6 Discussion.....	199
8.7 Summary	205
CHAPTER 9: STUDY 3 – FIELD EXPERIMENT	207
9.1 Purpose.....	207
9.2 Research Design.....	208
9.2.1 Measures	208
9.2.1.1 Independent Variable	208
9.2.1.2 Dependent Variables.....	210

9.2.1.3 Other Variables	210
9.2.2 Experiment Pretest	210
9.2.3 Experiment Execution.....	211
9.2.3.1 Sample.....	211
9.2.3.2 Training.....	212
9.2.3.3 Pretest.....	212
9.2.3.4 Task.....	214
9.2.3.5 Posttest	215
9.3 Analysis	216
9.4 Results.....	217
9.4.1 Overview across Phases	217
9.4.1.1 Phase One.....	217
9.4.1.2 Phase Two.....	219
9.4.1.3 Phase Three.....	220
9.4.2 Chosen Ideas	223
9.4.2.1 Process 1 – Individuals	223
9.4.2.2 Process 2 – Individuals and Collaborative Filtering	224
9.4.2.3 Process 3 – Individuals and Group Support System	226
9.4.3 Choice Shift	228
9.4.3.1 Individual-Level Choice Shifts	229
9.4.3.2 Group-Level Choice Shifts	234
9.4.4 Process Satisfaction and Decision Confidence	237
9.4.5 Participant Comments – Process and Tool	239
9.4.6 Comparison to Other Studies	241
9.4.6.1 Features	242
9.4.6.2 Challenges.....	242
9.4.6.3 Suggested Improvements	242
9.5 Discussion.....	243
9.6 Summary	248
CHAPTER 10: DISCUSSION AND CONCLUSION	250
10.1 Discussion.....	250
10.1.1 Research Questions.....	250
10.1.2 Key Findings.....	252

10.1.3 Design-Science Cycle	254
10.1.3.1 IS Research	255
10.1.3.2 Environment.....	257
10.1.3.3 Knowledge Base	257
10.1.3.4 Updated Design Theory	257
10.2 Contributions	258
10.2.1 Design Theory.....	258
10.2.2 Behavioural Theory	259
10.2.3 Practice.....	260
10.2.4 Calls for Additional Research.....	261
10.3 Strengths	263
10.4 Limitations	263
10.5 Future Research	264
10.6 Conclusions.....	265
BIBLIOGRAPHY	267
APPENDIX A: SAMPLE SIZE.....	287
APPENDIX B: TREATMENT CONDITIONS	289
APPENDIX C: THREATS TO VALIDITY.....	292
APPENDIX D: CARD SORT	296
APPENDIX E: QUESTIONNAIRES USED IN THE LAB EXPERIMENT	308
APPENDIX F: LAB EXPERIMENT – SCRIPT	332
APPENDIX G: LAB EXPERIMENT – TASK	337
APPENDIX H: PRELIMINARY ANALYSIS.....	338
APPENDIX I: CORRELATIONS	347
APPENDIX J: EDITING RULES	353
APPENDIX K: EXPLORATORY FACTOR ANALYSIS	354
APPENDIX L: MANCOVA.....	356
APPENDIX M: EXPERIMENTAL CONTROLS	388
APPENDIX N: QUALITATIVE DATA.....	397
APPENDIX O: IDEA EVALUATION	401
APPENDIX P: CASE STUDY – INTERVIEW GUIDES	423

APPENDIX Q: CASE STUDY – CODES	428
APPENDIX R: CASE STUDY – RESULTS	432
APPENDIX S: FIELD EXPERIMENT – TRAINING SCRIPT	436
APPENDIX T: QUESTIONNAIRES USED IN THE FIELD EXPERIMENT	446
APPENDIX U: FIELD EXPERIMENT – EXPERIMENT SCRIPT	464
APPENDIX V: FIELD EXPERIMENT – TASK.....	469

LIST OF FIGURES

Figure 2.1 – Knowledge Discovery and Similar Concepts.....	21
Figure 2.2 – Organizational Learning Process.....	34
Figure 2.3 – Integrated Learning Model.....	36
Figure 3.1 – Proposed Extensions to the 4I Model and Tool Support for Individual Learning Processes.....	53
Figure 4.1 – NKD Design Theory.....	71
Figure 4.2 – Athens: Iterative Clustering Approach.....	72
Figure 4.3 – Design Theory Process and the NKD Design Theory.....	78
Figure 5.1 – Updated Organizational Learning Process Model.....	80
Figure 9.1 – Field Experiment Phases and Processes.....	215
Figure 10.1 – Build-Evaluate Cycles for NKD Design Theory.....	256
Figure K.1 – Factor Plot: Knowledge-Discovery Type Items.....	354
Figure K.2 – Factor Plot: Learning Mode Items.....	355

LIST OF TABLES

Table 2.1 – Summary: Knowledge Discovery and Similar Concepts.....	20
Table 2.2 – Summary of Text- and Web-mining Techniques.....	24
Table 2.3 – Summary: Information Foraging and Similar Concepts.....	38
Table 3.1 – Learning Modes and Supporting Technologies.....	47
Table 3.2 – General Characteristics of Knowledge-Discovery Tools.....	48
Table 3.3 – Tool Descriptions.....	48
Table 3.4 – External Knowledge-Discovery Theoretical Framework.....	54
Table 4.1 – Design Theory Exemplars.....	60
Table 4.2 – Design Principles to Address Problem 1.....	65
Table 4.3 – Design Principles to Address Problem 2.....	67
Table 4.4 – Design Principles to Address Problem 3.....	69
Table 4.5 – NKD Design Theory Principles and Athens.....	74
Table 4.6 – Criteria Used to Assess Design Theory.....	75
Table 4.7 – General Characteristics of Knowledge-Discovery Tools.....	76
Table 5.1 – Decision Factors in Opportunity Evaluation.....	83
Table 6.1 – Summary of Propositions.....	93
Table 6.2 – Summary of Treatment Conditions.....	100
Table 6.3 – Organization Selection Criteria.....	101
Table 6.4 – Summary of Treatment Conditions.....	102
Table 6.5 – Criteria Used to Assess Experimental Design.....	103
Table 6.6 – Criteria Used to Assess Case Study Design.....	105
Table 7.1 – Summary of Covariate Variables.....	112
Table 7.2 – New Measurement Development Procedure.....	113
Table 7.3 – Summary of Pretest Activities.....	120
Table 7.4 – Summary of Variables.....	124
Table 7.5 – Factor Correlation Matrix: Knowledge-Discovery Types.....	129
Table 7.6 – Pattern Matrix: Knowledge-Discovery Types.....	130
Table 7.7 – Pattern Matrix: Knowledge-Discovery Types (All Transformed).....	131
Table 7.8 – Alternative Factor Groupings for Knowledge-Discovery Type Composite Variables	132

Table 7.9 – Summary Cronbach’s Alpha: Knowledge-Discovery Types.....	132
Table 7.10 – Rotated Factor Matrix: Learning Process	133
Table 7.11 – Pattern Matrix: Learning Process.....	134
Table 7.12 – Summary Cronbach’s Alpha: Learning Process	134
Table 7.13 – Summary Cronbach’s Alpha by Tool: Learning Process.....	135
Table 7.14 – Rotated Factor Matrix: Learning Mode	136
Table 7.15 – Rotated Factor Matrix: Learning Mode (3 Factor)	136
Table 7.16 – Summary Cronbach’s Alpha: Learning Mode.....	137
Table 7.17 – Pattern Matrix: All Variables.....	138
Table 7.18 – Pattern Matrix: All Variables (8 Factors)	140
Table 7.19 – Summary Cronbach’s Alpha: New Factors	141
Table 7.20 – Summary Cronbach’s Alpha: Original Factors.....	142
Table 7.21 – Rotated Factor Matrix: PIIT, GCSE, WSE.....	143
Table 7.22 – Summary Cronbach’s Alpha: PIIT, GCSE, WSE.....	144
Table 7.23 – Multivariate Tests of Significance	146
Table 7.24 – Support for Theory 1 Research Hypotheses	146
Table 7.25 – Support for Theory 2 Research Hypotheses	147
Table 7.26 – Support for Theory 3 Research Hypotheses	147
Table 7.27 – Differences between MANCOVAs	149
Table 7.28 – Qualitative Support: Task Manipulation Check.....	149
Table 7.29 – Qualitative Support: Tool Challenges.....	150
Table 7.30 – Support for Theory 3 – Knowledge Discovery Types	151
Table 7.31 – Support for Theory 3 – Learning Modes	152
Table 7.32 – Support for Theory 3 – Learning Processes.....	154
Table 7.33 – Results Summary	155
Table 7.34 – Features – Themes and Sample Comments	158
Table 7.35 – NKD Design Principles – Sample of Supporting Comments	158
Table 7.36 – Novel Knowledge – Supporting Comments	159
Table 7.37 – Challenges – Themes and Sample Comments	161
Table 7.38 – Suggested Improvements – Themes and Sample Comments	163
Table 8.1 – Summary of Participants.....	172
Table 8.2 – Summary: Case-Study Visit Schedule.....	174
Table 8.3 – Summary of Search Subtypes	182

Table 8.4 – Summary of Intuiting Themes	191
Table 8.5 – Summary of Interpreting Themes	193
Table 9.1 – Group Profile	211
Table 9.2 – Instructions for Idea Rating	218
Table 9.3 – Preliminary Decision Model – Phase 3.....	220
Table 9.4 – Final Decision Model – Phase 3	222
Table 9.5 – Process 1 Top Ideas – Results.....	223
Table 9.6 – Collaborative Filtering Idea Ranking.....	225
Table 9.7 – Process 2 Top Ideas – Results.....	226
Table 9.8 – Process 3 Top Ideas – Results.....	227
Table 9.9 – Summary of Top Ideas across Processes	227
Table 9.10 – Evaluation of Differences between Processes – Radicalness and Benefit.....	228
Table 9.11 – Overall Ratings: Individual Choice Shift between Phases 1 and 2.....	229
Table 9.12 – Radical Ratings: Individual Choice Shift between Phases 1 and 2.....	230
Table 9.13 – Benefit Ratings: Individual Choice Shift between Phases 1 and 2.....	230
Table 9.14 – Overall Ratings: Individual Choice Shift between Phases 2 and 3.....	231
Table 9.15 – Radical Ratings: Individual Choice Shift between Phases 2 and 3.....	232
Table 9.16 – Benefit Ratings: Individual Choice Shift between Phases 2 and 3	232
Table 9.17 – Differences between Choice Shifts – Individual Average Shift	233
Table 9.18 – Group-Level Choice Shift.....	235
Table 9.19 – Differences between Group Choice Shifts.....	236
Table 9.20 – Evaluation of Differences between Processes – Satisfaction and Confidence.....	237
Table B.1 – Summary of Treatment Conditions and Experimental Task	289
Table C.1 – Threats to Validity – Summary	292
Table C.2 – Construct Validity	293
Table C.3 – Statistical Conclusion Validity.....	294
Table C.4 – Internal Validity	295
Table D.1 – Round 1: Mapping of Categories Identified by Judges to Intended Constructs.....	296
Table D.2 – Round 1: Reworded Items.....	297
Table D.3 – Round 2: Mapping of Categories Identified by Judges to Intended Constructs.....	298
Table D.4 – Round 2: Reworded Items.....	298
Table D.5 – Round 3: Item Issues.....	299
Table D.6 – Round 4: Item Issues.....	300

Table D.7 – Round 5: Item Issues.....	301
Table D.8 – Summary of Items Dropped and Retained.....	302
Table D.9 – Benefit Items Dropped and Retained.....	303
Table D.10 – Radicalness Items Dropped and Retained.....	304
Table D.11 – Knowledge Discovery Items Dropped and Retained.....	305
Table D.12 – Learning Mode Items Dropped and Retained.....	306
Table D.13 – Learning Process Items Dropped and Retained.....	307
Table H.1 – Missing Value Analysis.....	338
Table H.2 – Univariate Outlier Analysis – Pre-transformation.....	340
Table H.3 – Normality Analysis: Athens.....	341
Table H.4 – Normality Analysis: Google.....	342
Table H.5 – Normality Analysis: Vivisimo.....	343
Table H.6 – Transformation Summary.....	344
Table H.7 – Multivariate Outlier Analysis – Post-transformation.....	346
Table I.1 – Correlation Matrix: Dependent Variables.....	348
Table L.1 – Correlation Matrix: Dependent Variables and Covariates.....	356
Table L.2 – Correlation Matrix: Covariates.....	357
Table L.3 – Multivariate Tests of Significance – Full Analysis.....	358
Table L.4 – Tests of Between-Subjects Effects – Full Analysis.....	358
Table L.5 – Adjusted Estimates – Full Analysis.....	361
Table L.6 – Pairwise Comparisons – Full Analysis.....	362
Table L.7 – Multivariate Tests of Significance – Knowledge-Discovery Type.....	363
Table L.8 – Tests of Between-Subjects Effects – Knowledge-Discovery Type.....	364
Table L.9 – Adjusted Estimates – Knowledge-Discovery Type.....	366
Table L.10 – Pairwise Comparisons – Knowledge-Discovery Type.....	367
Table L.11 – Multivariate Tests of Significance – Learning Processes.....	368
Table L.12 – Tests of Between-Subjects Effects – Learning Processes.....	369
Table L.13 – Adjusted Estimates – Learning Processes.....	371
Table L.14 – Pairwise Comparisons – Learning Processes.....	372
Table L.15 – Multivariate Tests of Significance – Learning Modes.....	373
Table L.16 – Tests of Between-Subjects Effects – Learning Modes.....	374
Table L.17 – Adjusted Estimates – Learning Modes.....	376
Table L.18 – Pairwise Comparisons – Learning Modes.....	377

Table L.19 – Multivariate Tests of Significance – Alternative Full Analysis	379
Table L.20 – Tests of Between-Subjects Effects – Alternative Full Analysis.....	380
Table L.21 – Adjusted Estimates – Alternative Full Analysis.....	382
Table L.22 – Pairwise Comparisons – Alternative Full Analysis.....	383
Table L.23 – Support for Theory 2 Research Hypotheses (Alternative Composites).....	385
Table L.24 – Support for Theory 3 Research Hypotheses (Alternative Composites).....	385
Table L.25 – Tests of Between-Subjects Effects – MP3 Knowledge	386
Table L.26 – Adjusted Estimates – MP3 Knowledge	386
Table L.27 – Pairwise Comparisons – MP3 Knowledge	387
Table M.1 – Multivariate Tests of Significance – Manipulation Check.....	388
Table M.2 – Tests of Between-Subjects Effects – Manipulation Check	388
Table M.3 – Estimates – Manipulation Check.....	389
Table M.4 – Pairwise Comparisons – Manipulation Check	389
Table M.5 – Manipulation Check – Tool Characteristics	390
Table M.6 – Multivariate Tests of Significance – Satisfaction, Training, Instructions	394
Table M.7 – Tests of Between-Subjects Effects – Satisfaction, Training, Instructions.....	394
Table M.8 – Estimates – Satisfaction, Training, Instructions	395
Table M.9 – Pairwise Comparisons – Satisfaction, Training, Instructions.....	395
Table M.10 – Differences in Satisfaction Between Tool Groups	396
Table N.1 – Features.....	397
Table N.2 – Challenges.....	398
Table N.3 – Suggested Improvements	399
Table O.1 – Summary of Benefit and Radical Items	401
Table O.2 – Missing Values and N/A Entries by Question	401
Table O.3 – Missing Values and N/A Entries by Rater	402
Table O.4 – Reliability Summary by Question Type.....	402
Table O.5 – Intraclass Correlation Coefficient (Ben1)	403
Table O.6 – Inter-Item Correlation Matrix (Ben1)	403
Table O.7 – Reliability Statistics (Ben1)	403
Table O.8 – Intraclass Correlation Coefficient (Ben3)	404
Table O.9 – Inter-Item Correlation Matrix (Ben3)	404
Table O.10 – Reliability Statistics (Ben3)	404
Table O.11 – Intraclass Correlation Coefficient (Ben5)	405

Table O.12 – Inter-Item Correlation Matrix (Ben5)	405
Table O.13 – Reliability Statistics (Ben5)	405
Table O.14 – Intraclass Correlation Coefficient (Rad1)	406
Table O.15 – Inter-Item Correlation Matrix (Rad1)	406
Table O.16 – Reliability Statistics (Rad1)	406
Table O.17 – Intraclass Correlation Coefficient (Rad2)	407
Table O.18 – Inter-Item Correlation Matrix (Rad2)	407
Table O.19 – Reliability Statistics (Rad2)	407
Table O.20 – Intraclass Correlation Coefficient (Rad4)	408
Table O.21 – Inter-Item Correlation Matrix (Rad4)	408
Table O.22 – Reliability Statistics (Rad4)	408
Table O.23 – Tests of Between-Subjects Effects – Idea Scores	409
Table O.24 – Estimates – Idea Scores	410
Table O.25 – Tests of Between-Subjects Effects – Idea Scores: Rater 1.....	411
Table O.26 – Estimates – Idea Scores: Rater 1.....	411
Table O.27 – Pairwise Comparisons – Idea Scores: Rater 1.....	412
Table O.28 – Tests of Between-Subjects Effects – Idea Scores: Rater 2.....	413
Table O.29 – Estimates – Idea Scores: Rater 2.....	414
Table O.30 – Tests of Between-Subjects Effects – Idea Scores: Rater 3.....	414
Table O.31 – Estimates – Idea Scores: Rater 3.....	415
Table O.32 – Pairwise Comparisons – Idea Scores: Rater 3.....	416
Table O.33 – Tests of Between-Subjects Effects – Idea Scores: Rater 4.....	417
Table O.34 – Estimates – Idea Scores: Rater 4.....	417
Table O.35 – Tests of Between-Subjects Effects – Idea Scores: Rater 5.....	418
Table O.36 – Estimates – Idea Scores: Rater 5.....	419
Table O.37 – Idea Composite Score Creation Techniques	420
Table O.38 – Multivariate Tests of Significance – Idea Composite Scores	420
Table O.39 – Tests of Between-Subjects Effects – Idea Composite Scores	421
Table O.40 – Estimates – Idea Composite Scores	422
Table Q.1 – Codes Used for Interviews (Group and Individual).....	428
Table Q.2 – Codes Used for both Interview and Tool-Use Observation	429
Table Q.3 – Codes Used for Tool-Use Observation	430
Table R.1 – Case Study and Lab Experiment – Comparison of Results.....	432

CHAPTER 1 INTRODUCTION

1.1 Novel Knowledge – A Case Example

Don Swanson, an information scientist, explores the medical literature to uncover novel knowledge. In his first investigation, Swanson sought to find novel knowledge about Raynaud's syndrome – a condition that results in intermittent restriction of blood flow to fingers and toes, triggered by cold or emotional stimuli (Gordon and Lindsay, 1996; Swanson, 1986; Swanson, 1990). At the time of his investigation, a cure for the condition had not been found. Although Swanson did not know what to look for specifically, his review of the literature resulted in the discovery of novel knowledge – a medical intervention for Raynaud's syndrome discovered by looking for indirect connections across several different bodies of literature. From his review of the Raynaud's literature, Swanson was able to make a connection between Raynaud's syndrome and blood viscosity. Upon review of the blood-viscosity literature, he found that dietary fish oil lowers blood viscosity. This led to the novel hypothesis that fish oil may be a useful dietary supplement to help decrease the blood viscosity in humans and therefore alleviate symptoms of Raynaud's syndrome (Swanson, 1986).

1.2 Introduction to Novel-Knowledge Discovery

Novel knowledge, defined as knowledge that is not currently known to the organization, interesting, relevant, but indirectly connected and therefore difficult to find, can be of significant strategic importance to the organization (Vats and Skillicorn, 2004a; Vats and Skillicorn, 2004b). Scanning the external environment and developing interpretations in order to “know” the environment is an important organizational activity (Daft and Weick, 1984). Some view the gathering of information about the environment as a basic requirement

for organizational survival (Daft and Weick, 1984). Others suggest it promotes growth (Penrose, 1959) and improved competitiveness (Daft and Weick, 1984). Further, the ability to understand and interpret information about the external environment enables entrepreneurial action, defined as the behaviours required for exploiting superior insights about the market relative to competitors (Sambamurthy et al., 2003; Zaheer and Zaheer, 1997). “Intuitively, it makes sense that the firm that knows more about its customers, products, technologies, markets, and their linkages should perform better” (Zack, 1999 p. 126). Organizations seek different types of knowledge in their external environments; novel knowledge is proposed as one of these types. Novel knowledge can guide organizations towards strategic opportunities that they are not aware of, but that are worth further investigation, for example dietary fish oil in the case above. This can provide an organization with a strategic advantage as the first-mover in learning (Zack, 2005). In addition, the ability to discover and learn novel knowledge, which depends on the organization’s absorptive capacity (Cohen and Levinthal, 1990), can enhance an organization’s learning capability and, as a result, provide a strategic advantage through “learning superiority” (Zack, 2002; Zack, 2005). Further, discovering knowledge that is novel to the organization can result in innovation (Majchrzak et al., 2004).

Sources of information about the environment can be classified as internal or external to the organization, and as personal or impersonal (Daft and Weick, 1984; Kourteli, 2005). The benefits and challenges vary with each type of source (Kourteli, 2005). The World Wide Web is a rich *external* source of information about the environment and is used extensively by organizations (Chung et al., 2005). However, the sheer volume and the dynamic nature of information on the Web creates a problem known as information overload: locating information is difficult because of the large amount of irrelevant information that must be sifted through (Bowman et al., 1994; Chung et al., 2005). There are a number of different tools and technologies that can help organizations find and make sense of information on the

Web, as well as direct the organization's attention to consequential information (Chung et al., 2005; Rao, 2004). Unfortunately, organizations face several challenges discovering novel knowledge on the Web, for example:

1. Difficulty locating novel knowledge in the vast amounts of information available online (e.g. Chung et al., 2005).
2. Difficulty recognizing novel knowledge, once located, as relevant and significant (e.g. Schulz, 2001).
3. Difficulty learning new concepts that are outside individual and organizational mental models (e.g. Hargadon and Douglas, 2001).

A novel finding may be strategically important to the organization, but if it is outside an individual's or organization's mental models, norms and underlying assumptions, novel findings may be interpreted as irrelevant.

Novel knowledge is a relative, subjective and socially-constructed concept. Knowledge may be perceived as novel to one individual or organization, but not another. The discovery of novel knowledge begins at the individual level and can move to the group and organizational levels if a shared understanding is developed. However, knowledge may be novel to an individual, but as it moves through the group and organizational levels, it may not be deemed novel.

In my dissertation, I examine how IT can support individuals, and in turn their organizations, in learning about and knowing their external environment on the Web. Specifically, I examine novel-knowledge discovery in the context of the multi-level organizational learning process. I begin this examination by developing a theoretical framework in which I identify the sets of tool characteristics, collectively referred to as levels, which are proposed to support the discovery of different types of knowledge, modes of learning and learning processes. As part of this theoretical framework, I propose extensions to the 4I organizational learning process model (Crossan et al., 1999) specific to searching and learning on the Web. Understanding which types of tools are useful in different discovery

and learning contexts has implications for learning effectiveness and may help firms understand how to “manage” their learning (Zack, 2005). Tools to support the discovery of novel knowledge, and the modification of mental models, are rare compared to tools supporting the discovery of other types of knowledge, for example, incremental knowledge that fine tunes existing mental models, and broader knowledge that expands existing mental models (Norman, 1982; Vandebosch and Higgins, 1996). Accordingly, I develop a novel-knowledge discovery (NKD) design theory – based on organizational and individual learning theories – that addresses an important and relevant problem for organizations: overcoming the challenges associated with novel-knowledge discovery. An instantiation of the design theory, a novel-knowledge discovery tool, is developed and tested and compared to tools at the other two levels of knowledge discovery. In addition, different processes involved in using a novel-knowledge discovery tool at the group level are examined. The results demonstrate how novel-knowledge discovery tools can support organizational learning.

This dissertation is organized as follows. I begin with a review of the knowledge discovery and organizational learning literatures in Chapter 2, including a more detailed description of the novel-knowledge construct. In Chapters 3, 4 and 5, I describe two research questions and the theoretical models proposed to address them: Chapter 3 describes the proposed external knowledge-discovery theoretical framework, including the proposed extensions to the 4I organizational learning process model (Crossan et al., 1999); Chapter 4 describes the proposed NKD design theory; and Chapter 5 discusses how different processes of NKD tool use – including both the generation of novel ideas and the evaluation and selection of an idea to pursue – affect the outcomes of tool use for the organization. In Chapter 6, I describe the research design used to examine the propositions identified in Chapters 3 through 5. In Chapters 7 through 9, I discuss in detail the methods used and results of three studies included in this research design: Study 1 – lab experiment; Study 2 – case

study; and, Study 3 – field experiment. I discuss the integrated findings from these three studies, as well as the contributions and limitations of this research in Chapter 10.

CHAPTER 2 LITERATURE REVIEW

In order to understand and examine external knowledge discovery and more specifically novel-knowledge discovery, the literatures related to novelty, knowledge discovery and organizational learning were reviewed.

2.1 Knowledge

Knowledge has been defined in a number of different ways (see Alavi and Leidner, 2001), based on a number of philosophical perspectives (Alavi and Leidner, 2001; Orlikowski, 2002). Traditionally, knowledge has been defined as justified true belief (Nonaka, 1994). Knowledge has also been distinguished from data and information: data represents facts; information represents organized data or data in context; knowledge represents information with direction, which enables action and decisions (Beccerra-Fernandez et al., 2004). There are different types of knowledge, each requiring a different management approach (Alavi and Leidner, 2001; Beccerra-Fernandez et al., 2004). In the literature, a differentiation between two types of knowledge – explicit knowledge and tacit knowledge – is often made (Alavi and Leidner, 2001; Nonaka, 1994; Polanyi, 1966). Explicit knowledge, also referred to as codified knowledge, is articulable in spoken or written form. Conversely, tacit knowledge is highly personal and embodied knowledge (Collins, 1993; Polanyi, 1966) that is not articulable and includes mental models, know-how, craft and skills (Nonaka, 1994). An individual gains or generates tacit knowledge through his/her experiences, or “learning by doing”. Explicit knowledge is articulable and can be exchanged in the form of information. Cultural knowledge, a third type of knowledge, refers to the assumptions, values and beliefs held by the individual or group (Choo, 2002). Other types of knowledge include declarative knowledge (know about), procedural knowledge (know how), causal knowledge (know why), and relational knowledge (know with) (Alavi and Leidner, 2001; Zack, 2001).

There have been numerous philosophical debates and perspectives regarding the definition and nature of knowledge. It is outside the scope of this study to review these philosophical debates and perspectives. However, of these different perspectives on knowledge, three views of knowledge seem particularly relevant in the context of Web searching and learning – object, state of mind and capability. As an *object*, knowledge is viewed as something to locate and manipulate (Alavi and Leidner, 2001). Once located, knowledge can affect an individual's *state of mind* in terms of his or her understanding and mental models (Alavi and Leidner, 2001). Lastly, the *capability* view of knowledge suggests that this change in understanding can increase the range of potential behaviors (Alavi and Leidner, 2001; Huber, 1991; Kim, 1993). Accordingly, the following definition of knowledge encompasses these three views: knowledge represents information with direction, which enables action and decisions (Beccerra-Fernandez et al., 2004) and increases the capacity for effective action (Huber, 1991; Kim, 1993).

2.2 Novelty

Novelty is a broad term that has been defined as the “appearance of something that is intrinsically new, i.e., something that is not contained within the environmental data and that is not a logical consequence of the unfolding of the previously existing state” (Encinar and Muñoz, 2006 p. 256; Schumpeter, [1932] 2005). The extent of newness, and to whom it is new, varies across definitions of novelty. Absolute novelty refers to an object or idea that is universally considered new (West and Anderson, 1996). When novelty is defined in this way, an entity is viewed as either novel or not (Encinar and Muñoz, 2006; Schumpeter, [1932] 2005; Styhre, 2006; Visser and Boschma, 2004). In contrast, relative novelty refers to an object or idea that is new relative to the unit of analysis (Zaltman et al., 1973). Thus, an idea could be considered novel to a particular organization, but not novel to another organization. Although these definitions treat the concept of novelty as a binary measure – novel *versus* not novel – other definitions treat novelty as discrete levels or a continuous measure. For example, innovations have been

measured by the degree of novelty, viewing the most novel innovations as radical and less novel ones as incremental (Encinar and Muñoz, 2006; Majchrzak et al., 2004). In addition, the concept of radical innovation has been broken down further, measuring the degree of radicalness, or novelty (Majchrzak et al., 2004). Innovations have also been measured by the overall level of *perceived* novelty (West and Anderson, 1996). Thus, the novelty label is relative, subjective and socially constructed (Schweizer, 2006).

Two processes related to novelty have been proposed (Witt, 2003): 1) the emergence of novelty, and 2) the dissemination of the innovation. The first process – the emergence of novelty – involves both the mental generation of ideas about new possibilities for action, as well as acting upon those possibilities to create an innovation. The emergence of novelty is related to entrepreneurial opportunities, which has been viewed from two perspectives: discovery and enactment (Dutta and Crossan, 2005). The discovery perspective would view a novelty as something in the environment that has not yet been identified, whereas the enactment perspective suggests that a novelty is enacted or created (Dutta and Crossan, 2005).

The concept of novelty can be differentiated from innovation in terms of learning – a cognitive or behavioural development (Fiol and Lyles, 1985). “Novelty is usually seen as emerging from a newly discovered possibility for action which, once taken, is called an innovation. However, any attempt to innovate is likely to trigger, and be accompanied by, learning” (Witt, 2003 p. 13). Thus, novelty can be viewed as a cognitive development and innovation a behavioural development. As a result, examining the cognitive-level learning literature is important for understanding novel-knowledge discovery.

The second process related to novelty – the dissemination of the innovation – occurs once the innovation has been developed. The dissemination process is primarily concerned with the adoption and diffusion of the innovation (Witt, 2003). The focus in this thesis is on the cognitive development aspects of the emergence of novelty. The process of generating novel ideas and the related challenges are discussed in more detail below.

“Novelty is an amorphous concept. By definition, the informational content, the meaning and the properties of what newly emerges, cannot be anticipated” (Witt, 2003 p. 38). Thus, novelty is related to the concept of surprise, as either counter expected or unimagined (Weick, 2006). The emergent and unpredictable nature of novelty suggests that it is difficult to search for a particular novelty. *Ex ante* specification of conditions or properties that a novelty should possess are based on existing knowledge of a problem or goal. However, the novelty that is actually discovered often reframes the current understanding of the problem or goal, rendering the *ex ante* conditions obsolete. Furthermore, if the conditions and properties of a novelty could be pre-specified, it would not be a novel idea (Witt, 2003).

Despite the inability to predict or anticipate novelty, the newly generated ideas are related to what the individual already knows to some degree. The process of generating novel ideas involves creating new mental concepts through a process of recombination. This process involves recombining existing cognitive concepts stored in memory, in some cases with external stimuli and new perceptions, resulting in the development of a new cognitive pattern and understanding (Barnett, 1953; Hargadon, 2002; Hargadon and Sutton, 1997; Witt, 2003). This recombination process is, therefore, constrained or bounded by what the individual already knows and the information “to which the agent’s environment draws his or her attention” (Witt, 2003 p. 39). Novel ideas that are completely unrelated to existing knowledge are rare.

Individuals will tend to produce more new ideas of possibilities for action than can be reasonably acted upon. Thus, individuals need to evaluate the set of ideas generated and choose those ideas upon which to act (Bontis et al., 2002; Witt, 2003). Ideas can be evaluated upon criteria such relevance and reliability. Relevance is an important criterion. Ideas that appear to be relevant, but may not be completely understood, may require additional research in order to fully evaluate the cost and benefits of the idea. The assessment of ideas is based on the individual’s existing knowledge base and subjective interpretation. Different individuals will not only

generate different novel ideas, but they will assess them differently too (Amabile, 1983; Amabile, 1988; Csikszentmihalyi, 1996; Witt, 2003).

2.3 Novel Knowledge

The term novel in this study is applied to knowledge and treated as a relative, subjective, socially-constructed and continuous measure. Novel knowledge represents knowledge that is new relative to the unit of analysis. Novel knowledge, even when known by another organization or industry, can lead to radical innovation when integrated with an organization's existing knowledge and applied to new problems or products (Hargadon, 2002; Hargadon and Sutton, 1997; Majchrzak et al., 2004; Zack, 2005).

Discovering novel knowledge begins at the individual level and moves to the group and organizational levels. The multi-level organizational learning process, discussed in more detail later, helps explain the flow of novel knowledge through these levels and the related challenges. Knowledge may be novel to an individual, but as it moves through the group and organizational levels, it may not be deemed novel. As a socially-constructed concept, knowledge is labeled as novel by the organization when members of the organization develop a shared understanding and perception of it as novel.

Novel knowledge should be interesting, arousing the attention of the organization, or the individual conducting the search¹. Interesting, in the knowledge-discovery context, has been defined as non-trivial, implicit, previously unknown and potentially useful (Fayyad et al., 1996). In addition, novel knowledge should be actionable. Uninteresting-and-new knowledge or knowledge that isn't actionable would likely be ignored or quickly discarded by the individual and organization.

Novel knowledge should be relevant to the organization in order for it to be useful and valuable. Thus, novel knowledge must be related, in some way, to what the organization does or

¹ From the American Heritage Dictionary, fourth edition.

knows. In concordance with the definition of novelty described above, novel knowledge is not directly connected to what the organization does or knows. The indirect connection of novel knowledge makes it difficult to find using conventional means (for example, using a common search tool such as Google). Thus, as both relevant and indirectly connected, novel knowledge is contextually appropriate; the indirect connection can be interpreted and understood in the organization's context.

Examples of novel knowledge help bring clarity to this amorphous concept. In Chapter 1, the example of the Raynaud's syndrome case was presented. In this case, Raynaud's syndrome was the focal-subject area with which the individual – Don Swanson – was familiar. Don Swanson was interested in finding something novel about Raynaud's syndrome in order to find a cure for the condition or some means of alleviating its symptoms – the context. Don's work surfaced a direct connection to blood viscosity, which led to the identification of the indirect connection to dietary fish oil. This indirectly related connection, which was contextually appropriate, spurred Don to develop the hypothesis – actionable knowledge – that dietary fish oil could alleviate the symptoms of Raynaud's syndrome. This hypothesis, which represents novel-knowledge discovery, moved to the group and organizational levels where medical professionals validated his hypothesis. This novel discovery turned into an innovation, which is now widely used and accepted by the medical profession and Raynaud's syndrome sufferers.

Another example of novel knowledge comes from an experiment run on the Web, back-dated to "September 12, 2001" (Skillicorn and Vats, 2004). In this experiment, the focal-subject area was bin Laden and the al Qaeda terrorist group. During this experiment, a strong yet indirectly related connection was found between Islamic terrorism and terrorist groups who are either non-Arabs or geographically distant from the Islamic terrorist groups (for example, Indonesian terrorists). This information could have been useful to US intelligence organizations in the aftermath of the September 11th attacks. The hypotheses generated as a result of these

indirect connections – resulting in actionable knowledge – may have helped shape the US’s strategy regarding military attacks and homeland security.

A more business-oriented example of novel knowledge comes from one of the experimental tasks used in this dissertation. In the lab experiment, students were asked to generate novel ways of applying their “hypothetical” company’s MP3 expertise. In this case, the focal-subject area was MP3 technology and industry. An example of an indirectly related connection is Wii², a video-game console with a wireless handheld-gaming device. A new product idea – actionable knowledge – linking MP3 technology with Wii gaming technology is an example of novel knowledge for this hypothetical organization.

Other examples of novel knowledge can be found in the innovation literature. For example, Hargadon (2002) describes a case in which a design team helped Reebok develop a response to Nike’s AirTM technology. The design team used their diverse knowledge to recognize indirect, yet contextually appropriate, connections between athletic shoes, inflatable splints, medical IV bags, pumps and valves. The novel knowledge that resulted was Reebok’s innovative PumpTM technology – use of inflatable air bladders in shoes to produce a form-fitting athletic shoe.

In another case, a project team from NASA searched for radically innovative ideas regarding instrument development for Mars exploration (Majchrzak et al., 2004). The NASA project team searched for existing solutions from other industries, which could be adapted to their context – measuring dirt devils on Mars. One of the resulting ideas involved recognizing an indirect, yet contextually appropriate, connection between space instruments for Mars and the British textile industry’s chair cover designs. The novel knowledge that resulted was an instrument design for measuring the electrostatic properties of dust on Mars.

² www.nintendo.com/wii

As these examples demonstrate, novel knowledge is related to the concept of recombinant innovations in which two existing and as-of-yet unrelated elements are combined to create something new (Hargadon, 2002; Hargadon and Sutton, 1997). For example, Raynaud's syndrome and fish oil were previously unrelated and were combined to create novel knowledge as an actionable idea.

2.4 Creativity

The concepts of creativity and novelty are intertwined. Some of the terms associated with creativity include problem solving, radical newness (Couger, 1990), novelty, value (Amabile et al., 1996; Barron and Harrington, 1981; Ghiselin, 1952; Mumford and Gustafson, 1988; Newell et al., 1962) and innovation (Mumford and Gustafson, 1988; Schweizer, 2006). Broadly, creativity had been defined as "any act, idea or product that changes an existing domain or that transforms an existing domain into a new one" (Csikszentmihalyi, 1996 p. 28). The following criteria were proposed to help identify creative thinking or products (Newell et al., 1962):

1. Thinking or products have novelty or value.
2. Thinking or products modify or reject current mental models.
3. Creative process requires persistence in order to achieve creative thinking or products.
4. Creative process includes the reformulation of the problem to be solved.

Across the varied definitions of creativity, three major factors emerge: novelty, utility or value, and recombination (Couger, 1990; Couger et al., 1993).

Four key components of creativity have been identified: person, process, product and press (environment) (Mooney, 1963). These components are discussed in more detail below.

2.4.1 Person

Individuals have creative potential that is viewed as either a permanent trait or a capability that can be learned (Barron and Harrington, 1981; Couger, 1990; Csikszentmihalyi,

1996; Kirton, 1976; Kirton, 2003). The creativity trait has been described as a cognitive style, which includes a continuum ranging from high adaption to high innovation (Kirton, 1976; Kirton, 2003). Creativity in the individual has been related to associative capabilities - the ability to explore and identify remote associations between elements, such as a words and concepts, in order to solve a problem (Couger, 1990; Mednick, 1962; Schweizer, 2006). Creativity has also been associated with divergent thinking – the ability to generate many different solutions to a given problem (Guilford, 1956). Although individuals differ in their creative ability and style, research has shown that individuals use less and less of their “native creativity” as they age because of various creative constraints in the environment, such as the emphasis on conformity in education and bureaucratic processes (Couger and Dengate, 1992).

2.4.2 Process

In contrast to the view that creativity is inherently possessed and expressed by the individual is the view that creativity can be facilitated and developed through a process. This creative process results in the generation of ideas (Amabile et al., 1996; Hender et al., 2002). Structured methods are advocated for “unstructuring” our thinking in order to generate ideas (VanGundy, 1988) – for example, brainstorming (Osborn, 1953). A variety of other techniques for enhancing creativity have been proposed, such as reformulating or reframing (Couger et al., 1993). Some authors have advocated a phased approach to the creative process (Csikszentmihalyi, 1996; Wallas, 1926), including an incubation period where individuals step away from the problem for a period of time in order to rest and unconsciously think about the problem. It has been proposed that at the end of the incubation period, creative ideas will emerge. Related to incubation is an approach referred to as a “change of venue” where the individual or group steps out of their typical environment into a new environment in order to stimulate the “creative juices” (Couger, 1990). Novel ideas are often spurred within unrelated contexts (Ghiselin, 1952). The creative process has also been described as an organic process, which also

requires some management (Ghiselin, 1952; Kanter, 1988). The initial processes involve “radical inventive activity” followed by refinement (Ghiselin, 1952).

2.4.3 Product

The output of the creative process is a product, either tangible or intangible. Although most research has focused on measuring an individual’s creative ability, the creativity of a product has also been investigated (Barron and Harrington, 1981; Couger and Dengate, 1992; Hender et al., 2002; Nagasundaram and Bostrom, 1995).

2.4.4 Press (Environment)

The fourth and final P of the 4 P’s model is “press”, which refers to the environment. The work environment and organizational culture in which creative work takes place is important. The work environment can either support or inhibit creativity (Amabile et al., 1996). In a related sense, there are a number of preconditions for creativity. These preconditions include the need for sufficient domain knowledge and creativity skills, as well as the motivation to be creative (Amabile, 1983).

A recent literature review focused on the influence of organizational culture on creativity and innovation (McLean, 2005). This review of the scholarly work of Amabile (e.g. 1983; 1988; 1996), Kanter (e.g. 1988) and Van de Ven, Angle and Poole (e.g. Angle, 1989), amongst others, synthesized past research and identified the most common organizational culture influences on creativity as: organizational encouragement, supervisory encouragement, work group encouragement, freedom and autonomy, and resources (McLean, 2005)³.

2.4.5 Groups and Creativity

In addition to research that examines the creativity of individuals, the study of groups and creativity has also been investigated. One type of group task that has received a lot of attention is

³ See McLean (2005) for a detailed review.

creativity tasks, which involve the generation of ideas or alternatives in groups. Creativity tasks often utilize techniques to increase the number of ideas generated and the novelty of those ideas. For example, brainstorming⁴, a frequently-used technique (e.g. Hender et al., 2002; McGrath, 1984; Potter and Balthazard, 2004), involves individuals within a group attempting to generate as many ideas as possible by thinking creatively and building on each other's ideas (Osborn, 1953). Although groups have the advantage of creating unique hybrid ideas due to interparticipant synergies, past research has found that individuals generate more, as well as more creative, ideas than groups as a result of process losses such as production blocking⁵ (Diehl and Stroebe, 1987; Diehl and Stroebe, 1991; Lamm and Trommsdorff, 1973; McGrath, 1984; Potter and Balthazard, 2004; Steiner, 1972). Recent research has recommended a multi-phase approach for idea generation by groups, including upfront phases where individuals prepare for and conduct brainstorming to generate ideas separately, followed by a phase where group members exchange and discuss ideas (Potter and Balthazard, 2004).

2.4.6 Relationship between Novelty, Creativity and Innovation

The concepts of novelty, creativity and innovation are highly related, crossing individual, group, organizational and societal-levels (McLean, 2005). As discussed above, novelty is one critical factor in the definition of creativity. Further, the outcome of the creative process is often an innovation. The relationship between novelty, creativity and innovation has been described in the Novelty Generation Model (NGM) (Schweizer, 2006). This model describes the novelty-generating process as consisting of three main constructs: novelty seeking, creativity, and innovative performance. Novelty seeking initiates the novelty-generating process and is an exploratory behaviour characterized by looking for and approaching unknown elements

⁴ A technique developed by Osborn (1953).

⁵ Production blocking refers to the inability of a group member to contribute an idea while another group member is speaking.

(Schweizer, 2006). This tendency towards novelty has been measured in the past using novelty-experiencing scales (Pearson, 1970).

The second component in the novelty-generating process is creativity, which is composed of two processes: novelty finding and novelty production. While novelty seeking involves the behaviour of looking for new things, novelty finding involves the ability to find new things – the ability to be creative. Thus, novelty finding occurs when the individual has the ability to think creatively, make intuitive leaps and develop unusual combinations of elements (Kirton, 1976; Schweizer, 2006). This creative ability has been evaluated in the past using tests such as the Remote Associates Test (RAT) (Mednick, 1962). Once the novelty has been found, the idea is put into action – novelty production.

The third component in the novelty generating process is innovative performance. Once the novel idea has been acted upon and developed into a product, the labeling of this product as an innovation and its performance as such depends on the evaluation and judgments of society/social environment. Thus, the innovative performance of this newly created product depends on “willingness and the ability to interact with the environment to get a product socially judged and recognized as novel” (Csikszentmihalyi, 1996; Schweizer, 2006 p. 168).

2.5 Knowledge Discovery

Knowledge discovery can be broadly defined as the development of new knowledge based on either the analysis of data and information, or the integration and reinterpretation of prior knowledge (Beccerra-Fernandez et al., 2004). External knowledge discovery refers to the discovery of knowledge outside the organization and therefore developing an understanding of the organization’s environment. Knowledge discovery is also a term that is used specifically in the IS and computing literature to describe a set of automatic and semi-automatic approaches for processing data and uncovering interesting, yet previously unknown, information, relationships and patterns (Grobelnik and Mladenic, 2005). Knowledge discovery is related to other constructs

such as knowledge creation, knowledge acquisition, knowledge sourcing, and knowledge exploration and exploitation. However, there are important differences, which are discussed below (see Table 2.1 and Figure 2.1 for a summary).

Knowledge creation involves developing new knowledge through the continual interaction and conversion between tacit and explicit knowledge (Choo, 2002; Nonaka, 1994). Over time, this interaction results in the amplification and spiral flow of knowledge, which moves from the individual to the group and organizational levels (Nonaka, 1994).

The four modes of knowledge creation are externalization, internalization, combination and socialization. These four modes involve converting explicit knowledge to tacit knowledge and *vice versa*, combining existing explicit knowledge to create new explicit knowledge, and converting tacit knowledge to new tacit knowledge (Nonaka, 1994). Most of these knowledge-creation modes involve social interactions, shared experiences, articulation and discussion of ideas within the organization, making them quite distinct from knowledge discovery. In contrast, the combination mode creates new explicit knowledge by “merging, categorizing, reclassifying, and synthesizing existing explicit knowledge” (Alavi and Leidner, 2001 p 116). Knowledge discovery is similar to the combination mode of knowledge creation, and could be considered to be a method of knowledge creation through combination. However, knowledge discovery is also focused on exploration and analysis processes, which knowledge combination does not appear to include.

Knowledge acquisition, often used synonymously with knowledge creation, refers to the process of obtaining knowledge (Huber, 1991). Part of the organizational learning process as described by Huber (1991), knowledge acquisition includes several sub-processes, such as congenital learning, experiential learning, vicarious learning, grafting and searching and noticing. Knowledge discovery is a method of knowledge acquisition, falling into the sub-process of searching and noticing. Examples of this sub-process include scanning – browsing without a specific question to answer – and focused search – looking for specific information (Huber, 1991;

Vandenbosch and Higgins, 1996; Vandenbosch and Huff, 1997). Knowledge discovery is distinguished by its focus on analysis, not exclusively on locating.

Knowledge sourcing is defined as “the extent to which an individual accesses other employees’ expertise, experience, insights, and opinions” (Gray and Meister, 2004 p. 821). Knowledge sourcing methods describe the mechanisms used to support and facilitate this access, such as knowledge repositories, meetings, etc. “Knowledge sourcing is an indirect learning behaviour where individuals gain access to others’ understanding of the work environment, mostly through language-based interactions” (Gray and Meister, 2004 p. 822). Three types of knowledge sourcing have been identified:

1. Dyadic knowledge sourcing – “a single knowledge seeker engaging in dialogue with an individual source” (Gray and Meister, 2004 p. 823).
2. Published knowledge sourcing – where many knowledge seekers access the knowledge of one source via codification. Thus, this type of knowledge sourcing does not include dialogue and is a one-way process.
3. Group knowledge sharing – “knowledge exchange in a setting containing multiple individuals” (Gray and Meister, 2004 p. 823).

Knowledge sourcing, like knowledge discovery, is focused on seeking knowledge that one does not currently have. The major difference is that knowledge sourcing is explicitly focused on acquiring knowledge from other individuals. Thus, knowledge sourcing is explicitly focused on expertise-based data. The data sources accessed in knowledge discovery can either contain factual data or expertise-based data. Further, knowledge discovery develops new knowledge through the analysis, integration or reinterpretation of existing content in vast data sources such as the Web. Although knowledge discovery may result in the discovery of a person, group or publication that leads to knowledge sourcing, it is not the purpose of knowledge discovery.

Exploration in the context of knowledge and learning refers to the activities that import new knowledge into the firm. Exploitation, on the other hand, refers to activities where existing knowledge is leveraged or applied to new uses (Bierly and Chakrabarti, 1996; Crossan and

Berdrow, 2003; Crossan et al., 1999; March, 1991; Zack, 1999). Knowledge discovery is a process of exploration and has been discussed in the literature as such (e.g. Chung et al., 2005; March, 1991). Knowledge discovery can be thought of as a specific form of exploration that involves significant analysis, integration and interpretation. Knowledge discovery is distinct from exploitation – it is focused on discovering new knowledge, whereas exploitation involves leveraging existing knowledge.

Table 2.1 – Summary: Knowledge Discovery and Similar Concepts

Concept	Similarities	Differences
Knowledge Creation	Similar to the combination mode of knowledge creation where new explicit knowledge is created by “merging, categorizing, reclassifying, and synthesizing existing explicit knowledge” (Alavi & Leidner 2001, p. 116).	The other three knowledge creation modes involve social interactions, shared experiences, articulation and discussion of ideas within the organization, which are not part of knowledge discovery. Knowledge discovery can also be distinguished by its focus on analysis and pattern recognition.
Knowledge Acquisition	Knowledge discovery is a method of knowledge acquisition, falling into the sub-process of searching and noticing.	Knowledge discovery can be distinguished by its focus on analysis, not exclusively on locating.
Knowledge Sourcing	Both focused on seeking knowledge that one does not currently have.	Knowledge sourcing is explicitly focused on acquiring knowledge from other individuals and focuses on expertise-based data. Knowledge discovery focuses on both factual and expertise-based data. Knowledge discovery develops knowledge through the analysis, integration or reinterpretation of existing content in vast data sources such as the Web.
Exploration	Knowledge discovery is a process of exploration.	The key difference is the analysis aspect of knowledge discovery.
Exploitation	None.	Knowledge discovery is focused on discovering new knowledge, whereas exploitation involves leveraging existing knowledge.

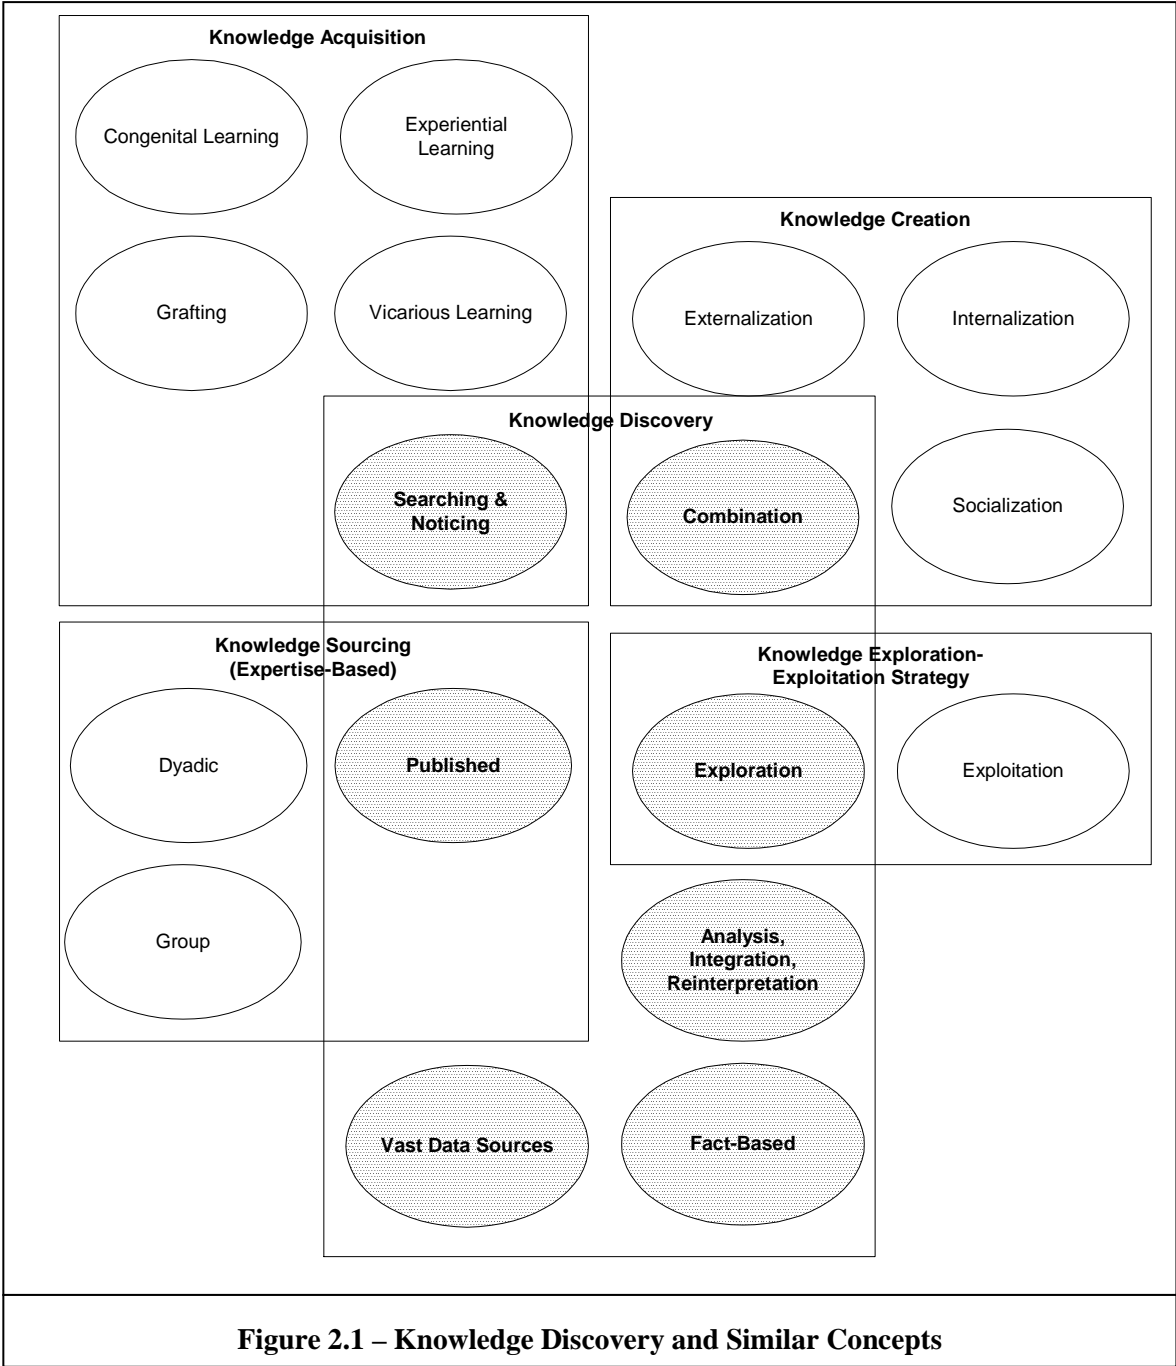


Figure 2.1 – Knowledge Discovery and Similar Concepts

2.6 Tools for Knowledge Discovery on the Web

The Web, while one of the top sources of information for businesses (Chung et al., 2005), presents a number of knowledge-discovery challenges, for example information overload. Accordingly, the Web is the context for the external knowledge discovery explored in this research. Two classes of tools for knowledge discovery on the Web are discussed below: data mining and search engines.

2.6.1 Data Mining

Recently, the term “data mining” has been used to describe the techniques used to process data and uncover interesting information, relationships and patterns; the term “knowledge discovery” is used to describe the overall process of defining the problem, searching, filtering, evaluating and interpreting these relationships and patterns (Fayyad et al., 1996; Tan et al., 2006). Thus, data mining refers to the collection of techniques used in the knowledge-discovery process.

There are a number of approaches included under the umbrella term of data-mining techniques, for example data mining, text mining and Web mining. Under the umbrella term of data-mining techniques, data mining refers to knowledge discovery that is focused on processing data in structured databases and unstructured data-sets. Since Web content, as one of the top sources of business information (Chung et al., 2005), is largely unstructured and text-based, this study is focused on Web- and text-mining techniques that process unstructured text, not on techniques to discover patterns in unstructured and structured data-sets and databases.

Text mining, an approach for discovering interesting information in unstructured text (Beccerra-Fernandez et al., 2004), uses information retrieval and artificial-intelligence techniques (Chen, 2001). Information retrieval involves text-manipulation and text-retrieval techniques (Grobelnik and Mladenic, 2005). Artificial-intelligence techniques focus on modeling and emulating human intelligence in computer systems (Chen, 2001). Since over eighty percent of information is stored as text, text mining is viewed as a potentially valuable tool for organizations

(Chen, 2001). Text can be analyzed at three different levels: lexical, syntactical and semantic. Lexical analysis focuses on individual words; syntactical analysis looks at the structure of text within statements; semantic analysis focuses on developing an understanding of the entire text (Grobelnik and Mladenic, 2005).

Historically, Web searches were lexical in nature. Semantic analysis has gained momentum in text mining as a result of the Semantic Web development – an initiative started by the World Wide Web Consortium to create a technological framework and universal medium for the exchange of data such that data can be shared and processed by automated tools as well as by people (Anonymous, 2005). The Semantic Web extends the current Web by attaching well-defined meaning to information through the use of ontologies (Berners-Lee et al., 2001). However, one of the challenges with ontologies is that they become dated and potentially obsolete over time (Grobelnik and Mladenic, 2005; Lytras, 2005; Rao, 2004). Text mining and semantic analysis are both related to content analysis, which also looks at deriving an understanding from textual documents (Rao, 2004).

Web mining refers to the analysis of Web-based documents using text-mining techniques⁶ (Chung et al., 2005). Types of Web mining include Web-content mining and Web-structure mining. Web-content mining analyzes Web pages to discover what the Web page is about and to discover knowledge, whereas Web-structure mining looks at the structure of links between a collection of Web pages (Beccerra-Fernandez et al., 2004). There are a number of different technologies and approaches used in Web mining, each suited to different types of applications. These technologies and approaches have been classified into four categories: linguistic analysis, co-occurrence analysis, clustering and categorization, and visualization (Beccerra-Fernandez et al., 2004; Chen, 2001) (see Table 2.2 for a summary).

⁶ In some cases, Web mining looks at the analysis of Web-usage patterns, also referred to as clickstream analysis (Beccerra-Fernandez et al. 2004)

Table 2.2 – Summary of Text- and Web-mining Techniques

Technique	Description
Linguistic Analysis	Techniques to identify key concepts in text documents. Examples: natural-language processing, information-extraction technologies, stemming, morphological analysis, Boolean, proximity, range, stoplists, semantic analysis and fuzzy search (Beccerra-Fernandez et al., 2004; Chen, 2001).
Co-occurrence Analysis	Utilizes indexing techniques to calculate the importance of terms in a document. Examples: Term-frequency-inverse-document-frequency, link analysis and similarity functions (Beccerra-Fernandez et al., 2004).
Clustering and Categorization	Clustering is a technique to classify “objects”, such as documents, into meaningful groups. Automatic categorization classifies “objects” into predefined groups within a taxonomy based on content similarity (Chung et al., 2005; Grobelnik and Mladenic, 2005; Rao, 2004).
Visualization	Techniques to display data graphically to help uncover patterns and relationships that can be visually detected, and to assist with exploration. Example: knowledge maps (Chung et al., 2005).

2.6.1.1 Linguistic Analysis

Linguistic analysis uses techniques, such as natural-language processing⁷, to identify key concepts in text documents (Beccerra-Fernandez et al., 2004; Chen, 2001). Various levels of analysis can be used in linguistic analysis: lexical, syntactical and semantic. Natural-language processing focuses on lexical analysis. Information-extraction technologies focus on developing an understanding of a document (Rao, 2004). This type of technology extracts what “objects” the document is referring to and what is being said about each object. Government, pharmaceutical and publishing organizations are utilizing this type of technology (Rao, 2004). Other linguistic techniques include stemming, morphological analysis, Boolean, proximity, range, stoplists, semantic analysis and fuzzy search.

2.6.1.2 Co-occurrence Analysis

Co-occurrence analysis utilizes indexing techniques to calculate the importance of terms in a document (Beccerra-Fernandez et al., 2004). The “term-frequency-inverse-document-frequency” function is a commonly used indexing technique. Term-frequency measures how often a term occurs in a document. Inverse-document-frequency looks across a number of documents to determine the specificity or uniqueness of a term. A term has a higher index value when it appears in fewer documents. “Term-frequency-inverse-document-frequency” multiplies

⁷ Analysis of natural language.

the term-frequency by the inverse-document-frequency to identify terms used frequently in one document, but infrequently across a number of documents (Beccerra-Fernandez et al., 2004). Other techniques include link analysis and similarity functions (Beccerra-Fernandez et al., 2004).

2.6.1.3 Clustering and Categorization:

Clustering is a technique to classify “objects”, such as documents, into meaningful groups (Chung et al., 2005). Different clustering techniques use different definitions of “meaningful”, but in general, a measurement of similarity is used to group objects into clusters. Clustering techniques may also leverage heuristic techniques, such as genetic algorithms, taboo search, and simulated annealing (Blum and Roli, 2003), to determine how many clusters to create.

Automatic categorization is similar to clustering. However, in the case of automatic categorization the groups are predefined (Grobelnik and Mladenic, 2005; Rao, 2004). Documents are assigned to a category within the existing taxonomy based on content similarity. One of the problems with automatic categorization is the “dead taxonomy” issue, where the predefined taxonomy becomes out of date or obsolete over time (Rao, 2004).

2.6.1.4 Visualization

Visualization techniques display data graphically to help uncover patterns and relationships that can be visually detected, and to assist with exploration. Visualization allows data to be represented in multiple dimensions. A “knowledge-map”, which displays relationships between objects such as Web pages, is one visualization output example (Chung et al., 2005).

2.6.2 Web Search Engines

Search engines, tools commonly used to locate knowledge on the Web, enable users to query across a domain of content, based on a set of search criteria. Rather than searching the contents of the Web itself, the search engine compiles an index of Web content and searches across this index. A search engine typically produces a long list of results, listing the Web sites that relate to the search criteria, sorted by relevance (Beccerra-Fernandez et al., 2004). Relevance

criteria, which can vary across search engines, are crucial because they influence the type of knowledge discovered by the individual and learning that results. Examples of different relevance algorithms are weighted keyword scores, and page ranking, which are based on the number and quality of Web sites that link to the target Web site (Chen, 2001). The page ranking approach, which is utilized by Google, uses a “global” computation to calculate the relevance of a page, based on how the rest of the world assesses that page’s relevance (i.e., by linking to that page themselves) (Chen, 2001). However, global relevance does not necessarily mean that the page is the most interesting and relevant in this individual’s particular context. More advanced search tools organize and display results in clusters and knowledge maps⁸ in order to help the user understand the deeper structure of the results. In addition, the Semantic Web can support the discovery of knowledge on the Web that is more relevant to the individual’s context.

Data mining and knowledge discovery, although focused on the automated techniques and technologies for uncovering hidden relationships and patterns, are exploratory and learning processes. Despite the focus on eliminating the need for human intervention in knowledge discovery, the process of exploration and discovery does include humans, and is embedded within human, group and organizational contexts. If the individuals within an organization do not understand the new knowledge discovered, the organization will not learn about this new knowledge. Furthermore, if the individuals learn new concepts but cannot transfer this knowledge and learning into the organization, the organization will not learn the new knowledge discovered. Thus, in addition to the automated techniques described above, knowledge discovery is affected by individual, group and organizational learning processes. In order to understand how to design a tool to facilitate novel-knowledge discovery, it is important to examine these processes, which are discussed in the following section.

⁸ A knowledge map represents the results in a graphical map, with nodes (circles) as Web sites and lines linking the nodes together to represent common key words between nodes (Chung et al., 2005).

2.7 Organizational Learning

Organizational learning has been defined in a number of different ways: focusing on cognitive development, behavioural development, or both (Bontis et al., 2002; Crossan and Berdrow, 2003; Crossan et al., 1999; Edmondson, 2002; Fiol and Lyles, 1985; Hargadon and Fanelli, 2002; Weick and Roberts, 1993). Organizational learning has also been defined as an increase in the range of potential behaviours (Huber, 1991) or capacity to take effective action (Kim, 1993). An alternative interpretation is that the organization has learned a new behaviour, but has not necessarily performed the behaviour. Organizational learning focusing on cognitive development looks at the development of insights, knowledge, and associations between past actions, the results of those actions and future actions (Fiol and Lyles, 1985). Organizational learning focusing on behavioural development looks at action and changes in organizational behaviours and structures (Fiol and Lyles, 1985; Simon, 1969). While some organizational learning definitions combine both behavioural changes with changes in insight (Crossan and Berdrow, 2003; Crossan et al., 1999; Hedberg, 1981; Simon, 1969; Weick and Roberts, 1993), others distinguish these two types of changes as adaptation and learning respectively (Fiol and Lyles, 1985). In other words, organizational learning involves the development of insights, knowledge and associations; adaptation involves making incremental behavioural and structural changes as the environment changes (Fiol and Lyles, 1985). Cognitive developments do not necessarily lead to changes in behaviour; changes in behaviour do not necessarily mean that new insights have also been gained (Fiol and Lyles, 1985). However, there may be a reciprocal relationship between cognition and action (Crossan and Berdrow, 2003; Crossan et al., 1999; Edmondson, 2002; Hargadon and Fanelli, 2002; Kim, 1993). This research views organizational learning as cognitive developments since the focus of this thesis is on discovering new knowledge, not on the new behaviours that result from this knowledge discovery.

In the following sections, the literature on learning, distinguishing between cognitive developments and the learning process to achieve those cognitive developments, is reviewed. In addition, the impact of tools on the learning process and learning outcomes is discussed.

2.7.1 Mental Models

The knowledge-discovery process is both enabled and constrained by mental models, which are frameworks to help simplify and organize information (Crossan et al., 1999; Day, 2002; Hedberg, 1981; Kim, 1993). Both individuals and organizations have mental models. Individuals have their own unique mental models; whereas the organization's mental model is a shared, negotiated understanding (Fiol and Lyles, 1985; Hedberg, 1981; Kim, 1993) or patterns of heedful interrelating (Weick and Roberts, 1993). Mental models have also been referred to as cognitive structures (Nystrom and Starbuck, 1984), conceptual frameworks (Levitt and March, 1988), schemas (Hargadon and Douglas, 2001), theories of action (Argyris and Schon, 1978; Hedberg, 1981), and frames (Barsalou, 1992; Orlikowski and Gash, 1991; Orlikowski and Gash, 1994). Collectively, these terms refer to knowledge structures that represent knowledge as a complex network of concepts with abstract attributes, values, relationships and rules (Barsalou, 1992). These knowledge structures, which I refer to as mental models, have also been conceptualized as a collection of routines and frameworks (Kim, 1993). Rather than containing individual pieces of data, mental models are aggregates of data that "prescribe a viewpoint or course of action" (Hedberg, 1981; Kim, 1993 p. 45). Although mental models are useful for guiding interpretation and making sense of environmental information, they affect what the individual and the organization searches for and sees in the environment (Hargadon and Fanelli, 2002; Kim, 1993; Orlikowski and Gash, 1991). Mental models, as simplifying frameworks used to order new information, tend to reinforce existing models and are often taken for granted and unarticulable (Argyris and Schon, 1978). As a result, mental models are difficult to change and

“sticky”⁹. Mental models are dangerous if they remain unexamined and unchanged because they can blind the individual and the organization to novel knowledge and potentially important opportunities or threats in the environment (Day, 2002; Hargadon and Fanelli, 2002; Weick, 2006).

Organizational learning theories provide insight into how organizations can overcome the challenge of “sticky” mental models. Organizational learning theories differentiate between single-loop learning and double-loop learning (Argyris and Schon, 1978; Fiol and Lyles, 1985)¹⁰. Single-loop learning, a low-level of learning, involves repetition of past successful behaviours, incremental learning and change to routines within existing rules, norms and frames of reference (Crossan et al., 1999; Fiol and Lyles, 1985; Huber, 1991). Thus, single-loop learning, a behaviour-related form of learning, is most closely related to what is currently known. Double-loop learning, a higher level of learning, involves surfacing, challenging and changing norms, routines, assumptions, and mental models that were previously inaccessible because they were either unknown or known, yet undiscussable (Argyris and Schon, 1978; Fiol and Lyles, 1985; Kim, 1993). Double-loop learning, a highly cognitive form of learning, also involves changing frames of reference and developing complex associations and causal relationships (Fiol and Lyles, 1985; Huber, 1991). Double-loop learning may allow organizations, constrained by their mental models, to change and discover novel knowledge.

2.7.2 Learning Modes

Learning mode describes the degree to which mental models are changed when new information is acquired (Hedberg, 1981; Norman, 1982; Vandenbosch and Higgins, 1996). Based on cognitive learning theories (e.g. Norman (1982), Piaget (1954) and Maier (1945)),

⁹ The term “sticky” is used here to refer to the difficulty in modifying mental models and is not meant to refer to the term used to describe the difficulty in transferring knowledge between areas within the organization (Szulanski, 1996).

¹⁰ Argyris and Schon (1978) also discuss a third type of learning – deuterio learning – which involves learning to learn and developing a learning capability.

Vandenbosch and Higgins (1996) proposed two modes of learning: mental model building and mental model maintenance. The first mode, mental model building, involves the creation of new mental models. When new concepts do not fit into existing mental models or disconfirm current models, mental models may be fundamentally changed in order to align with this new knowledge (Vandenbosch and Higgins, 1996). The second mode of learning, mental model maintenance, involves the confirmation of existing mental models. When new information fits into existing mental models, this information is added to and helps validate the existing mental model (Vandenbosch and Higgins, 1996).

In the organizational learning literature, additional learning modes have been discussed. Reinventive learning (Gnyawali and Stewart, 2003), radical learning (Edmondson, 2002), turnaround learning (Hedberg, 1981) and innovation (Gray and Meister, 2004) involve questioning and making fundamental or radical changes to existing mental models. Similar to reinventive learning, formative learning involves creating entirely new mental models in order to provide a new understanding of existing information (Gnyawali and Stewart, 2003). Reinventive, innovation, radical, turnaround and formative learning are most similar to the mental model building mode discussed above, as well as double-loop learning, which involves surfacing, challenging and changing mental models (Argyris and Schon, 1978). Adjustive learning (Gnyawali and Stewart, 2003), turnover learning (Hedberg, 1981) and adaptation (Gray and Meister, 2004) involve making incremental changes and refinements to existing mental models. Adjustive learning, turnover learning and adaptation are similar to Norman's (1982) tuning mode of learning, where mental models are slowly refined and updated over time. Lastly, operative learning (Gnyawali and Stewart, 2003), adjustment learning (Hedberg, 1981), incremental learning (Edmondson, 2002) and replication (Gray and Meister, 2004) involve validating and reinforcing existing mental models. Operative, adjustment, incremental, replication, and single-loop learning, which involves repetition of past successful behaviours and change to routines

within existing frames of reference (Crossan et al., 1999; Fiol and Lyles, 1985; Huber, 1991), is most similar to mental model maintenance.

Learning modes and the degree to which mental models change are related to imagination and creativity. Imagination involves creating and conceiving of something novel and the exponential growth of an idea. Imagination requires redefining what is known, breaking free from past experiences, and reframing (Weick, 2006). Thus, imagination requires mental model building. In contrast, fancy involves creating something new within the confines of existing mental models and incremental idea growth (Weick, 2006). Thus, fancy involves mental model maintenance.

The link between mental models and creativity can be found in brainstorming research. The cognitive processes involved in brainstorming, as well as techniques to stimulate certain cognitive processes, have been examined in past research (e.g. Hender et al., 2002; Nagasundaram and Dennis, 1993; Potter and Balthazard, 2004). This research suggests that during idea generation, individuals search through familiar categories and associations in their memory (Anderson, 1987; Anderson, 1992; Barsalou, 1983; Hintzman, 1988) and mental models (Piaget, 1954; Piaget and Inhelder, 1969). Brainstorming productivity can be improved by providing individuals with cues or stimulus to help with this memory search and the generation of ideas (Hoffman, 1959; Nagasundaram and Dennis, 1993; Potter and Balthazard, 2004). Further, the creativity of the ideas generated is affected by the degree to which the cues or stimuli are related to the problem (VanGundy, 1988) and individuals' existing mental models.

When individuals are provided with stimulus that is related to the problem, they tend to search narrowly, accessing familiar associations in their existing mental models. Examples of techniques that provide related stimulus include: 1) cause cueing (Potter and Balthazard, 2004), where individuals think of problem causes using their own "natural" categories (Barsalou, 1983), and 2) prompts, where individuals are asked multiple questions related to the problem (Hender et al., 2002). Using techniques involving related stimulus is less cognitively complex and, as a

result, more efficient than techniques involving unrelated stimulus because natural categories and familiar associations are triggered (Barsalou, 1983; Hender et al., 2002; Nagasundaram and Dennis, 1993; Potter and Balthazard, 2004). However, existing mental models constrain creativity and, thus, ideas tend to be lower in creativity (Hender et al., 2002; Nagasundaram and Bostrom, 1995). From a learning mode perspective, related stimulus triggers mental model maintenance (Vandenbosch and Higgins, 1996).

Alternatively, when individuals are provided with stimulus that is unrelated to the problem, the act of forcing a relationship back to the problem (VanGundy, 1988) requires the individual to make unfamiliar and remote (or uncommon) connections, reframing their existing mental models (Hender et al., 2002), and triggering mental model building (Vandenbosch and Higgins, 1996). Although this process is more cognitively complex and more time consuming, the result is highly creative ideas (Hender et al., 2002; Nagasundaram and Bostrom, 1995; Nagasundaram and Dennis, 1993).

The process of organizational learning, discussed in more detail below, has also been linked to imagination, creativity and radical innovation. Recent research has proposed that organizations in today's "hyperdynamic" environment need to incorporate creativity and radical innovation into their learning process in order to gain a sustainable competitive advantage (Wang and Ahmed, 2003).

2.7.3 Individual vs. Organizational Learning

The organizational learning literature describes how organizational learning is both related to and distinct from individual learning. Individual learning has been defined as an increase in an individual's capacity to take effective action (Kim, 1993). Although individual learning contributes to organizational learning (March, 1991), what an organization learns is not simply the sum of what individuals in the organization learn (Argote, 1999; Crossan et al., 1999; Fiol and Lyles, 1985; Hedberg, 1981; Kim, 1993). Organizational learning is clearly dependent

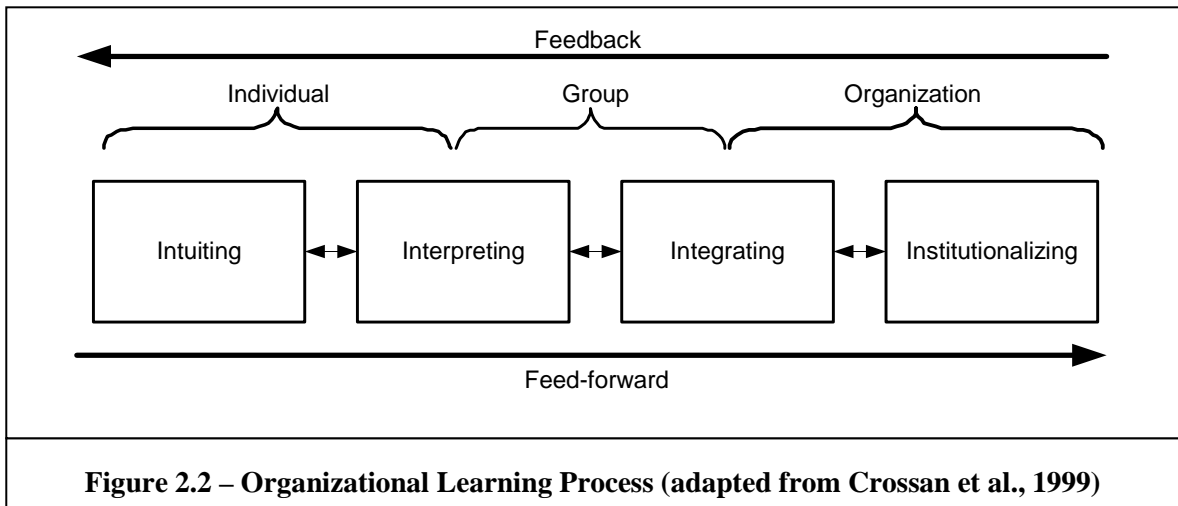
on individual learning (Argyris and Schon, 1978), but it is not dependent on what any specific individual learns (Kim, 1993). A transfer mechanism between individual learning and organizational learning is needed. Some have argued that the key distinction between organizational learning and individual learning is sharing: within the organization and over time, shared interpretations are developed and negotiated (Daft and Weick, 1984). Individual and organizational learning are connected to one another as part of the organizational learning process. This connection is explored in more detail in the organizational learning process section below.

In addition to individual learning contributing to organizational learning, organizational learning influences individual learning and behaviour. Organizations develop learning systems, which influence current and future members of the organization. These learning systems embed what the organization learns over time into systems, structures, strategy, routines, and norms (Bontis et al., 2002; Crossan et al., 1999; Fiol and Lyles, 1985; Hedberg, 1981; March, 1991; Stein and Zwass, 1995). Essentially, learning systems help develop the organization's memory, which preserves "behaviours, mental maps, norms, and values over time" (Hedberg, 1981 p. 6) such that past events and experience influence the organization's current activities (Stein and Zwass, 1995). The elements of an organization's memory that affect how the organization chooses to act, what it chooses to remember, and what it chooses to attend to, are referred to as the organization's "active" memory (Kim, 1993). An organization's memory and what it currently knows affects what information it searches for, as well as its perception and interpretation of the environment and its experiences (Huber, 1991).

2.7.4 Organizational Learning Process

The organizational learning process has been conceptualized as a multi-level dynamic process, including both feed-forward and feedback processes (Crossan et al., 1999). These processes span the individual, group and organizational levels, and include the following:

intuiting, interpreting, integrating and institutionalizing (referred to as the 4I model). Intuiting is an individual-level cognitive process, which involves pattern recognition and the development of insights into past patterns and future possibilities. At this point, ideas are not yet formed, but are subtle, fuzzy, initial inklings of interesting possibilities. Interpreting begins as an individual process and can move to a group process where ideas are formed, articulated, discussed and debated. The process of interpreting involves developing models for understanding, bringing out meaning, and developing shared mental models amongst managers (Crossan et al., 1999; Daft and Weick, 1984). The next process, integrating, is a group process where a shared understanding is translated into coordinated action. The last process, institutionalizing, occurs at the organization level when learning becomes embedded into routines, rules, procedures and infrastructures. Figure 2.2 depicts the organizational learning process (Crossan et al., 1999).



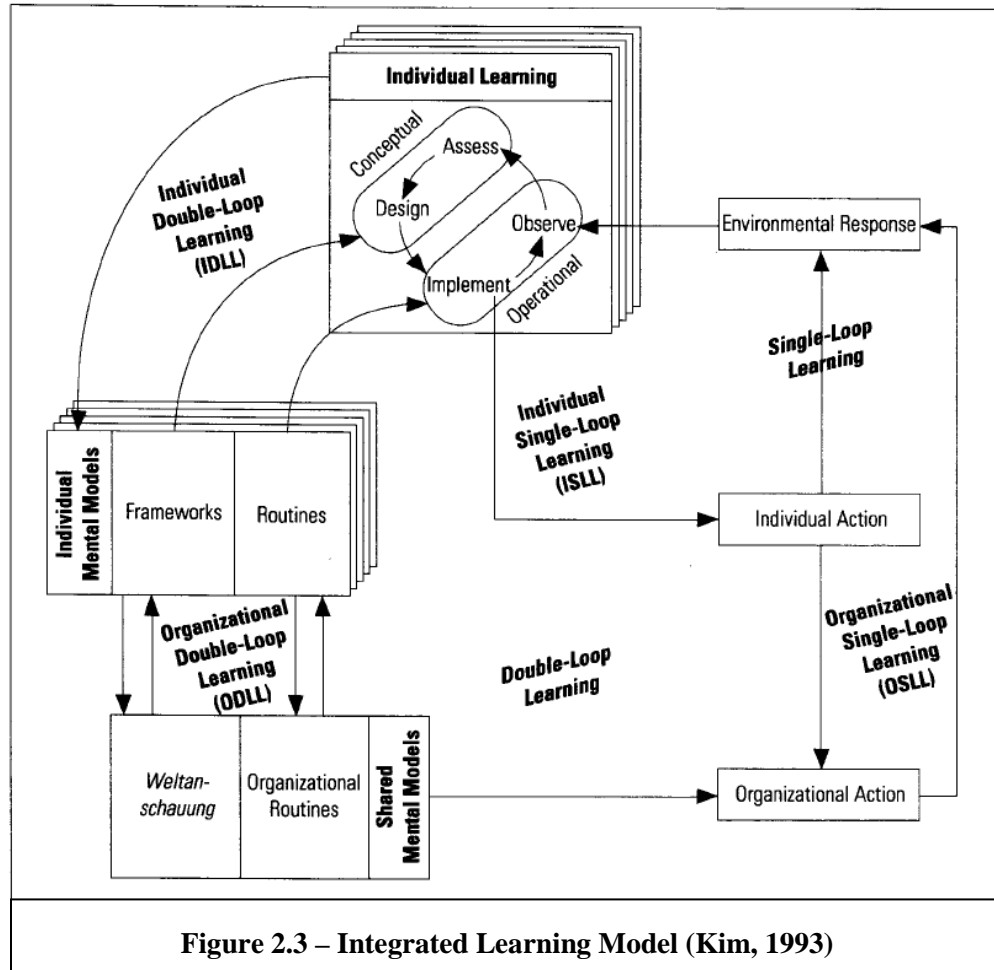
The feed-forward process of organizational learning, the focus here, is an exploratory process that translates individual and group insights into learning that is institutionalized via “organizational learning systems” (Crossan et al., 1999; March, 1991). Cognitively, the connection between individual and organizational learning is the transfer and absorption of individual mental models into the organization’s shared mental models (Kim, 1993). The group

processes of interpreting and integrating, where a shared, negotiated understanding is developed, mediate this process.

The feedback process exploits institutionalized learning, which guides individual and group thinking and action (Crossan et al., 1999). However, institutionalized routines, norms and mental models are slow to change and, thus, may constrain and hinder future learning and exploration (Crossan and Berdrow, 2003; Crossan et al., 1999; March, 1991).

As noted above, the group plays an important role in the transfer of individual learning to organizational learning. In the 4I model, the group serves as a mediator in the organizational learning process. Others view the “group” in a slightly different light. Edmondson (2002) views groups as the most important entity in the organizational learning process. Organizational learning is conceptualized as a “process of cascading team learning opportunities, independently carried out, but interdependent in their impact on company performance” (Edmondson, 2002 p. 144). Each group has different learning goals, different learning processes and, thus, learning is viewed as both local and variegated.

Alternatively, Kim (1993) conceptualizes the organizational learning process as a two-level process model, including individual and organizational learning (see Figure 2.3 below). Although the group is not included in this model, Kim (1993) suggests that the group can be viewed as a “mini-organization”, where individual members of the group contribute to the group’s shared mental models. Thus, the model can explain the group learning process in addition to the organizational learning process. The group can be viewed as a “collective individual” that contributes to the organization’s shared mental models. Kim’s (1993) treatment of the “group” in the organizational learning process suggests that the individual-to-group learning process is the same as, or highly similar to, the individual-to-organization learning process.



The group learning processes of interpreting and integrating, discussed above, are closely related to the research on group tasks and processes. McGrath integrated past group-task research – for example, Shaw (1973; 1976), McGrath and Altman (1966) and Steiner (1972), amongst others – to create the Group Task Circumplex, which identifies four main task groups or quadrants and a total of 8 group-task types. These tasks are further differentiated by the degree to which they are conceptual *versus* behavioural and involve conflict *versus* cooperation. Creativity and decision-making tasks, both conceptual task types, are the most relevant task groups within the scope of this research. Creativity tasks, included in the “generate” task quadrant, involve the generation of ideas or alternatives. Decision-making tasks, included in the “choosing” task

quadrant, involve choosing an alternative when there is no “correct” answer¹¹. Often referred to as preference tasks, decision-making tasks require reaching a consensus regarding the preferred alternative. Creativity and decision-making tasks both span the group interpreting and integrating learning processes.

Other organizational learning models have included processes such as scanning, learning (Daft and Weick, 1984), reconceptualizing the problem, deciding to search, search and evaluate, and development (Majchrzak et al., 2004). The scanning, search and evaluation processes can be conceptualized more broadly as information foraging, discussed in more detail below.

Information foraging refers to the process of seeking information, as well as assessing and choosing information sources based on the expected value of information and expected cost of locating and extracting this information (Pirolli and Card, 1995; Pirolli and Card, 1999). Thus, information foraging involves assessing the trade-off between the value of information gained *versus* the costs of performing the search activity. Information foraging does not focus on how to find more information, but on how to optimize the user’s time. Information foraging is based on the “optimal foraging theory” from biology and anthropology, which looks at food-foraging strategies in animals (Pirolli and Card, 1995; Pirolli and Card, 1999). Information foraging is an evolutionary process that involves identifying an initial search-space, browsing the search-space to identify the required information and, if required, reformulating the search terms to refine and narrow the search space (Hoare and Sorensen, 2005). The search strategy adapts to the constraints and affordances of the task at hand – the context of the knowledge sought. For example, search strategies for complex and important knowledge will differ from strategies to search for simple and “interest only” knowledge. The search strategy also adapts to the expected value of information and expected cost of searching.

¹¹ For example, in mathematical models, there may be a “correct” solution.

While information foraging is related to concepts such as information retrieval, scanning and focused search, it is a distinct construct (see Table 2.3). Information retrieval is a broad term that spans a continuum of retrieval behaviours from a specific, focused search to a broad, unfocused search (i.e. scanning) (Vandenbosch and Huff, 1997). While both information retrieval and foraging refer to an active acquisition of information, information foraging specifically describes an evolutionary and adaptive strategy for finding valuable information and minimizing the associated costs. It involves calculated moves in search space to optimize the chance of finding something of value. Information foraging can involve both focused search and scanning behaviours as part of its adaptive and evolutionary search strategy.

Table 2.3 – Summary: Information Foraging and Similar Concepts

Concept	Similarities	Differences
Information Retrieval	Both actively attempt to find information.	Information foraging is a more specific term that describes an evolutionary and adaptive strategy for finding valuable information and minimizing associated costs. It involves calculated moves in search space to optimize the chance of finding something of value.
Focused Search	Both involve specific, calculated and focused search behaviours.	Information foraging changes the search strategy in order to adapt to the task and search environment.
Scanning	Information foraging may begin with a broad scan of the search space in order to understand how to formulate a more specific focused search.	Information foraging changes the search strategy in order to adapt to the task and search environment.

2.7.5 Learning Novel Concepts

The literature reveals competing views regarding the most effective way to change mental models and learn new concepts. One view suggests that to learn something possessing a high degree of novelty and modify mental models, current mental models must either be separated from the learning process, or discarded altogether. An example of separation is “opportunistic learning”, where organizations purposely bypass existing standard procedures in order to seize an opportunity or learn new concepts (Kim, 1993). Some organizations adopt a “skunkworks” approach, an example of opportunistic learning, to overcome the rigidities of their institutional context and organizational mental models (Hargadon and Fanelli, 2002; Quinn,

1996; Tushman and O'Reilly III, 1999). A skunkworks is a relatively small organizational unit that has been physically, culturally and structurally separated from the rest of the organization. Separation from the rest of the organization frees the skunkworks from the constraints of the larger organization. Organizations use a skunkworks approach when innovation is needed, for product development or addressing difficult-to-solve problems, and when the organization's institutional context makes innovation difficult. Opportunistic learning suggests that separating learning from existing mental models is most effective for learning novel concepts.

Rather than requiring a separation from current mental models, some research suggests that discarding knowledge and mental models is necessary for learning novel concepts (e.g. Crossan and Berdrow, 2003; Hedberg, 1981; Nystrom and Starbuck, 1984; Wang and Ahmed, 2003). An example of discarding current mental models is removing top managers, a drastic move that is partly ritualistic, but may be necessary to unlearn and change mental models (Nystrom and Starbuck, 1984). Top managers are reluctant to change their beliefs and mental models, which strongly influence the organization's beliefs and mental models, and will tend to interpret most information as confirming the validity of their mental models. Thus, removing top managers and discarding mental models is sometimes the most effective way to learn and develop new mental models.

In contrast to separating and discarding current mental models, the second view argues that learning new concepts must be grounded in familiar knowledge and mental models (e.g. Cohen and Levinthal, 1990; Hargadon and Douglas, 2001). For example, innovation adoption and diffusion research suggests that for individuals to learn, understand, and adopt innovations, they must be linked to current mental models and familiar ideas (Hargadon and Douglas, 2001). In fact, innovations that do not tie to existing and familiar mental models may never be understood and adopted. Hargadon and Douglas (2001) suggest that the design elements of an innovation must enable individuals to develop their interpretation of the novelty by merging elements of what is familiar with elements of what is new. For example, the digital video recorder has

features that are similar to the VCR, a familiar tool, but other features that are novel. Thus, the innovation is strongly linked to what is understood and familiar, and attempts to extend the mental model from familiar and known elements to novel and thus unknown elements.

An alternative view, based on Majchrzak's work on knowledge reuse for innovation (Majchrzak et al., 2004), is that learning novel concepts may be enabled through reframing – conceptualizing existing mental models and knowledge in a new frame (Watzlawick et al., 1974). The change literature discusses two type of change: first-order and second-order change. These two types of change are closely related to the concepts of single- and double-loop learning respectively. Thus, first-order change refers to change within a given framework or system, whereas second order change refers to changing the framework or system itself (Watzlawick et al., 1974). Reframing is viewed as a mechanism for creating second-order change (Watzlawick et al., 1974). When reframing, one does not change the objective facts or the object, rather, one just classifies it or perceives it differently (Watzlawick et al., 1974). Reframing or reconceptualizing an existing problem, for example, can result in solutions that provide a discontinuous improvement for the organization (Kim, 1993), or a new goal to motivate the search for opportunities and ideas (Majchrzak et al., 2004). However, the reframing or new framework has to be related in some way to the person's existing knowledge or mental models, otherwise it will not be absorbed and understood (Watzlawick et al., 1974).

The ability to learn new concepts is also dependent on the individual's or organization's absorptive capacity (Cohen and Levinthal, 1990) and learning readiness. Learning readiness in adults¹² is influenced by the tasks and responsibilities associated with his or her evolving role in the organization (Knowles, 1970). The influence of role and task on learning novel concepts was evident in the results of the three studies included in this dissertation, discussed in Chapters 7 through 9.

¹² The antecedents of learning readiness differ between children and adults.

2.7.6 Learning and IT

Past research has explored how IT can support learning effectively. The executive support system (ESS) literature has looked at the modes of learning supported by ESS, depending on how the ESS is used and characteristics of the user. Vandebosch and Higgins (1996) explored the relationship between how the ESS was used by executives – for scanning or focused search – and the resulting mode of learning – mental model maintenance or mental model building. In a separate study, Vandebosch and Huff (1997) explored which system characteristics influenced scanning behaviour in an ESS. For example, the ability to integrate information from different sources and flexibility in how information can be used were found to be important system features that supported scanning behaviour (Vandebosch and Huff, 1997).

The organizational learning literature has explored the impact of different types of tools (e.g. groupware, knowledge repositories and communication tools such as email) on learning processes within the organization – exploitation and exploration (Kane and Alavi, 2005; Kane and Alavi, 2007). Using computational modeling, both speed of learning and variance of knowledge were found to differ across types of tools, thus supporting exploitive and explorative processes differently (Kane and Alavi, 2005; Kane and Alavi, 2007).

The virtual learning environment (VLE) literature examined the important design dimensions of a VLE, specifically the degree of learner control, to support learning effectiveness (Piccoli et al., 2001). Similarly, Leidner and Jarvenpaa (1995) explored the types of tools and technologies that best support different learning models (objectivist, constructivist, collaborative, cognitive information processing, and sociocultural)¹³ in the business school environment. However, further work is needed to understand the importance and suitability of different tool design elements in different learning environments and different learning models (Piccoli et al., 2001).

¹³ See Leidner and Jarvenpaa 1995 for a detailed discussion of these learning models.

2.8 Summary: Integrating Knowledge, Learning, and IT

The concepts discussed above – knowledge discovery, learning modes and processes, and tools – are interrelated. Learning modes describe the degree to which mental models – internal representations of knowledge – change. The learning process helps explain *how* mental models are changed and affect future learning. The 4I organizational learning process model (Crossan et al., 1999) is used in the theoretical framework developed in Chapter 3 because it describes how mental models are created at the individual, group and organizational levels and how these mental models affect interpretation and sense-making. Thus, it enables an integrative analysis of learning mode and process.

Finally, tools can support individuals and organizations in the discovery of knowledge – based on a “discovery” approach to novelty emergence. This discovery provides a stimulus for learning and initiates a learning process involving intuiting and interpreting, both processes of “enacting” (Dutta and Crossan, 2005). Learning is guided by existing mental models and results in updates to those mental models. This learning process can also be supported by tools.

In Chapter 3, the research questions regarding novel-knowledge discovery, learning and IT are described. In addition, a theoretical framework integrating these concepts is proposed. Chapters 4 and 5 further develop specific areas of the theoretical framework.

CHAPTER 3 THEORETICAL FRAMEWORK

3.1 Introduction

The main research questions to be addressed in this dissertation are as follows:

1. What design properties should a novel-knowledge discovery tool have in order to be effective?
2. How do tools oriented towards different levels of knowledge discovery affect the organizational learning process?

In order to address these research questions, a theoretical framework is proposed below, describing the relationship between tools, knowledge discovery and learning. In addition, a design theory for novel-knowledge discovery (NKD) is proposed in Chapter 4, followed by a series of propositions in Chapter 5 comparing different processes of using an NKD tool in terms of outcomes for organizational learning.

An external knowledge-discovery theoretical framework is proposed below describing a series of relationships between tool characteristics, which are grouped together in different levels, and the type of knowledge discovered, learning mode and learning processes supported.

3.2 External Knowledge-Discovery Theoretical Framework

The learning modes identified by Vandenbosch and Higgins (1996) – mental model maintenance and mental model building – encompass most of the learning modes described by Norman (1982), Maier (1945) and Piaget (1954). However, there is a third mode of learning described by Norman (1982) that is mentioned, but not explicitly included in their research model because it was not applicable in the ESS context being studied. The third mode of learning is tuning, a mode that involves the continual adjustment of knowledge within an existing mental model. Existing knowledge is refined, made more efficient, specialized and embedded into the mental model as one encounters more experience or data to refine and update the details of that

knowledge. Thus, existing mental models are refined and specialized expertise is developed (Norman, 1982). Thus, the following three modes of learning are included in this framework:

1. *Mental model maintenance*: confirmation and validation of existing mental models (Vandenbosch and Higgins, 1996).
2. *Mental model tuning*: refinement and updates to existing mental models (Norman, 1982).
3. *Mental model building*: creation of new mental models or fundamental changes to existing mental models (Vandenbosch and Higgins, 1996).

The amount of learning and change to existing mental models rises from mode 1 (mental model maintenance) to mode 3 (mental model building).

As explored in other system domains such as ESS (Vandenbosch and Higgins, 1996; Vandenbosch and Huff, 1997) and VLE (Piccoli et al., 2001), it is useful to understand what tool functionality, in general, supports different modes of learning on the Web. Dworman et al. (2000) discuss two types of questions that are relevant for searching in either structured (e.g. databases and indexed text-based search and retrieval systems) or unstructured (e.g. full-text documents) sources of data: 1) questions about trees, and 2) questions about the forest. Questions about trees are “record-oriented” questions (Dworman et al., 2000). In other words, these questions are about specific things. This relates to the “mental model maintenance” mode of learning since the answers to these specific questions help to confirm and validate existing knowledge. This is supported by past research linking *focused-search* behaviour to mental model maintenance (Vandenbosch and Higgins, 1996). From a technology perspective, tools that support searching for specific things (i.e. focused search) should be suited to “mental model maintenance”, for example standard search engines such as Google (www.google.com).

The questions about the forest are pattern-oriented questions (Dworman et al., 2000), in other words, finding patterns of relationships amongst words and concepts in the case of textual data. Finding new relationships between related concepts can help individuals refine and update existing mental models with new connections that fit into these mental models. Thus, questions

about patterns relate to the “mental model tuning” mode of learning. Tools that highlight patterns and display them to users should support this mode of learning, for example clustering tools such as Vivisimo (www.vivisimo.com), or knowledge-mapping tools such as Kartoo (www.kartoo.com).

The third learning mode, mental model building, may not involve a distinct question; rather, the question may be vague and unarticulable. Mental model building may involve looking and scanning for unknown connections. This is supported by past research linking *scanning* behaviour to mental model building (Vandenbosch and Higgins, 1996). Kimbrough refers to these types of unknown connections as “surprising associations” or low-probability relationships (Kimbrough, 2001). Norman (1982) suggests that structuring, a mode related to mental model building, involves sudden changes in knowledge and the restructuring of knowledge. These unknown relationships highlight connections that are indirect and linked through an intermediate connection (Gordon and Lindsay, 1996). Kimbrough refers to questions oriented towards unknown relationships as “Raynaud questions” after Swanson, the information scientist who discovered a medical intervention for Raynaud’s syndrome by looking for these indirect connections across the medical literature (Kimbrough, 2001). This led to the discovery that dietary fish oil supplements could help alleviate the problems associated with Raynaud’s disease, an indirect connection that was not found in any of the documents reviewed (Swanson, 1986; Swanson, 1990).

These unknown and indirect connections can lead to the discovery of novel knowledge. Tools that can help individuals find these unknown and indirect connections are most suitable for mental model building. However, tools of this nature are rare. There are some tools that can look across a document collection (for example the Core of Discovery (Kimbrough, 2001)); however, there are no tools that I am aware of that can perform this function across the Web. Table 3.1 summarizes the proposed connections between learning modes, types of questions addressed and supporting tools. The general tool characteristics proposed to support different learning modes

are summarized in Table 3.2. For level 1 and 2 tools, the characteristics were developed by generalizing the functionality within each family of tools proposed to support a specific learning mode. Since tools at level 3 are rare, a design theory for novel-knowledge discovery tools is developed in Chapter 4 and Table 3.2 updated accordingly. Table 3.3 provides a description of specific tools at each level.

Table 3.1 – Learning Modes and Supporting Technologies

Learning Mode	Description	Related Learning Theories	Questions and Answers	Supporting Technologies
Mental Model Maintenance (Vandenbosch and Higgins, 1996)	Confirms and validates existing mental models (Vandenbosch and Higgins, 1996).	Reproductive thought (Maier, 1945); assimilation (Piaget, 1954); accretion (Norman, 1982); operative learning (Gnyawali and Stewart, 2003); incremental learning (Edmondson, 2002); replication (Gray and Meister, 2004); single-loop learning (Argyris and Schon, 1978); adjustment learning (Hedberg, 1981).	Questions: Specific, record-oriented questions (Dworman et al., 2000; Huber, 1991; Vandenbosch and Higgins, 1996). Answers: Focused, specific knowledge.	Search engines, semantic web technologies. Tool Example: Google
Mental Model Tuning (Norman, 1982)	Refines and updates existing mental models; existing knowledge made more efficient, specialized (Norman, 1982).	Tuning (Norman, 1982); adjustive learning (Gnyawali and Stewart, 2003); adaptation (Gray and Meister, 2004); turnover learning (Hedberg, 1981).	Questions: General, pattern-oriented questions (Dworman et al., 2000). Answers: Patterns, relationships between related concepts.	Categorization and clustering tools, ontology development tools, semantic web technologies. Tool Example: Vivisimo
Mental Model Building (Vandenbosch and Higgins, 1996)	Creates new mental models (Vandenbosch and Higgins, 1996), involving sudden changes and restructuring of knowledge (Norman, 1982).	Productive thought (Maier, 1945); accommodation (Piaget, 1954); structuring (Norman, 1982); reinventive and formative learning (Gnyawali and Stewart, 2003); double-loop learning (Argyris and Schon, 1978); innovation (Gray and Meister, 2004); radical learning (Edmondson, 2002); turnaround learning (Hedberg, 1981).	Questions: Vague or unarticulable. Answer: Novel knowledge and “surprising associations” (Kimbrough, 2001).	Text mining tools. Tool Example: Unknown

Table 3.2 – General Characteristics of Knowledge-Discovery Tools

Tool Level	General Characteristics	Tool Examples
Level 1	<ol style="list-style-type: none"> 1. Searches across a domain of content using keywords entered by the user. (*) 2. Produces list of results matching search terms. (*) 3. Includes links to relevant content and meta-knowledge of that content. (*) 4. Sorts results by relevance / similarity to search terms (for example, weighted keyword scores, and page ranking). (*) 	Google, Yahoo, Altavista, Infoseek, Ask, Dogpile, Metacrawler ¹
Level 2	<ol style="list-style-type: none"> 1. Includes characteristics of level 1 tools (where noted *). 2. Groups results into clusters of similar content, thus identifying various dimensions of the search term and related concepts. 3. Allows user to drill down into clusters for more detailed results. 	Kartoo, Vivisimo, NorthernLight, Excalibur, Autonomy, Verity ²
Level 3	To be developed in Chapter 4.	Athens 1.0 (partially)

Notes: ¹ Both Dogpile and Metacrawler have some features that are consistent with level 2, such as asking the user "Are you looking for x?" and allowing the user to drill down into different search word combinations.
² Now part of Autonomy's product offerings.

Table 3.3 – Tool Descriptions

Tool Example	Description
Google (Level 1)	Google (www.Google.com), a popular search engine, allows users to query a domain of content based on a set of search criteria. Google produces a long list of results, listing the Web pages that correspond to the search criteria. The results are sorted by relevance, specifically by page ranking, which is based on the number and quality of Web pages that link to the target Web page (Chen, 2001).
Vivisimo (Level 2)	Vivisimo (www.Vivisimo.com) is a search and clustering tool that organizes and displays the results of a search in groups according to similarity of content, developing a taxonomy. Users can navigate and drill down into the taxonomy to find specific Web pages according to their interests.
Athens (Level 3)	Athens 2.0 is developed in this research, an extension of Athens 1.0 (Vats and Skillicorn, 2004a; Vats and Skillicorn, 2004b), and represents an instantiation of the NKD design theory. Athens 1.0 partially addresses the NKD design principles. Athens 2.0 will fulfill those design principles not addressed in the current version of Athens. Athens accepts several parameters, including at least two keywords and Web domain specification (for example, IBM.com or *.uk), and returns information from the Web domain that is indirectly related to that keyword. The tool uses an iterative clustering technique to discover novel knowledge.

Based on the connections between learning modes, types of questions addressed and supporting tools identified above, the proposed levels of knowledge discovery are: level 1, focused on discovering incremental and specific details about what is already known; level 2, focused on gaining an understanding of the multiple dimensions of a concept, and thus broader,

enriched knowledge; and level 3, focused on uncovering novel knowledge and surprising associations. Thus, I propose:

P1: Knowledge-discovery levels can be differentiated by the type of knowledge sought.

P2: Associated levels of knowledge-discovery tools exist.

More specifically, I propose the following:

P2a: Level 1 knowledge-discovery tools¹ most effectively discover incremental details about what is already known, in order to deepen one's understanding of a concept.

P2b: Level 2 knowledge-discovery tools² most effectively discover multiple dimensions of a concept and relationships with other concepts, in order to broaden one's understanding of a concept.

P2c: Level 3 knowledge-discovery tools³ most effectively discover novel knowledge and surprising associations.

P2d: Level 1 knowledge-discovery tools support mental model maintenance.

P2e: Level 2 knowledge-discovery tools support mental model tuning.

P2f: Level 3 knowledge-discovery tools support mental model building.

While each level of tool can be used to discover all types of knowledge and support all learning modes (for example, Google may be used to discover novel knowledge), it is proposed that each tool level is most *effective* at discovering and supporting these particular knowledge types and learning modes.

To extend these initial propositions, I propose that the three levels of tools also support individual learning processes differently and, thus, add a fourth level to the 4I (1999) learning process model: the machine-level. As discussed above, the two individual learning processes are intuiting and interpreting. In order to integrate different organizational learning process models, information foraging, a very relevant process in the context of searching and learning on the Web, is added as a fifth learning process to the 4I (1999) learning process model. Information foraging,

¹ Such as Google.

² Such as Vivisimo.

³ Such as Athens.

in the context of Web-searching, may support the individual learning process of *intuiting* – cognitive processes involving pattern recognition and developing insights – and *interpreting* – developing models for understanding and bringing out meaning (Crossan et al., 1999). Adding the machine-level and information foraging to the 4I (1999) learning process model integrates knowledge-discovery and learning processes.

The development of appropriate search terms is important for locating information on the Web. The development of search terms is an intuitive and inductive process, which starts with the intuiting process and feeds into the information-foraging process. When the initial search term(s) do not provide the desired results, the intuiting process feeds back into the information-foraging process for an additional iteration. Thus, the development of search terms connects the intuiting and information-foraging processes, forming an iterative loop (see Figure 3.1 below). The degree to which different levels of tools support each of these four learning processes – search term development, information foraging, intuiting and interpreting – is discussed below.

Level 1 tools, for example standard search engines, support the information foraging process for individuals by locating information and assessing its usefulness for the individual based on the search terms provided. The long list of results that are returned by most search engines creates information overload (Bowman et al., 1994; Chung et al., 2005) and does little to support the individual intuiting and interpreting processes. However, these processes are less important when individuals are looking for specific knowledge. Level 1 tools rely solely on the individual for the development of search terms, which requires intuition and induction.

Level 2 tools help to organize and display results in clusters and knowledge maps in order to help the user find relevant and interesting results based on the search terms entered. Level 2 tools support information foraging and partially support the interpretation process – assigning meaning to information. For example, a knowledge map can highlight clusters related to a topic and the relationships between these clusters, making it is easier for the individual to understand the topic, sub-topics and relationships between them. However, knowledge maps can

become so large that a user has difficulty understanding the scope of the map and difficulty finding what is most interesting. Thus, intuiting – developing insights and recognizing patterns – may be difficult. Level 2 tools also rely on the individual for the development of search terms.

The goal of level 3 tools is to help discover novel knowledge – knowledge that the individual or organization is not currently aware of, that is relevant, and indirectly connected to what the individual or organization knows (Vats and Skillicorn, 2004a; Vats and Skillicorn, 2004b). All knowledge discovered is new to the recipient; the important distinction is the degree of “newness” and surprise to the recipient. One of the distinguishing features of novel-knowledge discovery (NKD) tools, as differentiated from the other levels, is that the tool directs the individual’s attention to what they could do or learn next based on what is indirectly connected, yet related to the search terms provided (Vats and Skillicorn, 2004a) (see Table 3.2 for a summary of tool characteristics). This feature helps reduce the information overload problem experienced in the other levels. As discussed above, tools at levels 1 and 2 use relevance algorithms that tend to bury interesting results in long lists, or require significant browsing. Level 3 tools are proposed to point the individual’s attention to a smaller number of interesting and relevant ideas in the result set.

In addition to supporting information foraging, level 3 tools fully support the interpreting process and partially support the development of search terms. Since the user, by the definition of novel knowledge, does not know what he or she is looking for, level 3 tools should refine the search terms initially provided by the user or at least be less reliant on search-term precision. Further, in order to discover indirect connections that are related in a meaningful way to the initial search terms, level 3 tools need to be able to assess and choose the series of logical connections that will likely lead to meaningful results for the user, for example the connections between Raynaud’s syndrome, blood viscosity and fish oil in the case of Swanson’s discovery. This supports the interpreting process of sense-making and developing an understanding of the results. Level 3 tools also partially support intuiting – pattern recognition and the development of

insights – by directing an individual’s attention and insight to novel and potentially important knowledge.

Figure 3.1 depicts the proposed extensions to the 4I organizational learning process model (Crossan et al., 1999) and how tool levels 1, 2 and 3, respectively, support individual level learning processes differently. Accordingly, I propose the following:

P3a: Level 1 knowledge-discovery tools⁴ are able to support the information foraging process.

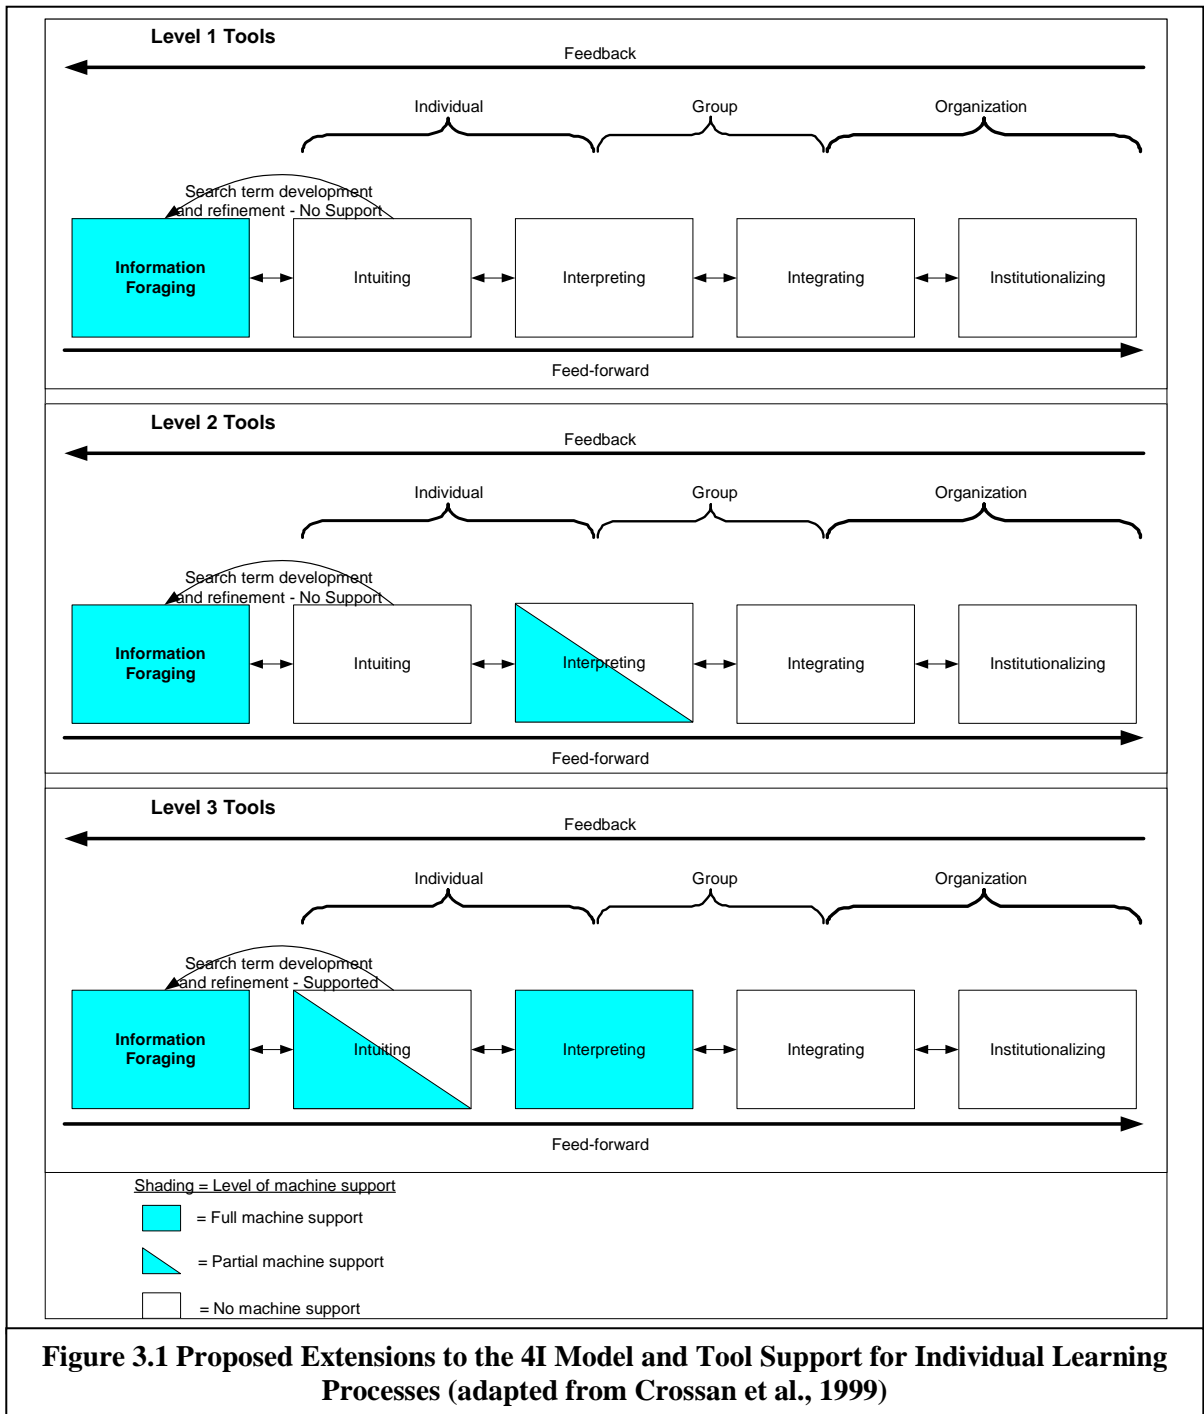
P3b: Level 2 knowledge-discovery tools⁵ are able to support the information foraging process and partially support the interpretation process.

P3c: Level 3 knowledge-discovery tools⁶ are able to support the information foraging and interpretation processes, and partially support the intuiting process, including search-term development.

⁴ Such as Google.

⁵ Such as Vivisimo.

⁶ Such as Athens.



Based on propositions 1 through 3c, the external knowledge-discovery theoretical framework is summarized below in Table 3.4. This framework describes how tool levels are

proposed to relate to knowledge-discovery type (i.e. the types of answers found), learning process and learning mode.

Table 3.4 – External Knowledge-Discovery Theoretical Framework

Tool Level	Knowledge Discovery Type	Learning Process Supported	Learning Mode
Level 1 Tools (for example, Google)	Specific, focused, “deep” knowledge	Information foraging	Mental model maintenance
Level 2 Tools (for example, Vivisimo)	Patterns, relationships, “broad” knowledge	Information foraging; Interpreting (partial)	Mental model tuning
Level 3 Tools (for example, Athens)	Surprising associations, “novel” knowledge	Information foraging; Intuiting (partial); Interpreting	Mental model building

While the theoretical framework identified above is parsimonious, there are some important caveats to note about this framework. The type of knowledge discovery or, rather, the type of question the individual is asking is context specific. In one context, an individual may be looking for a specific answer. In other situations, an individual involved in the innovation process may want to find novel knowledge. Further, the answers found may not align with the questions asked and associated tool level. An individual looking for a specific, focused answer using a level 1 tool may discover surprising associations and novel knowledge. Thus, it is possible to discover novel knowledge with a level 1 tool such as Google. However, it may require significantly more effort or serendipity to do so.

The type of knowledge discovered is also dependent on the individual’s existing knowledge and perceptions. Each individual has a unique set of mental models and experiences, which influences how the question is framed and search terms developed, as well as how the results are interpreted. For example, one individual might view the knowledge as novel and find interesting new opportunities as a result. Another individual might find the knowledge confirms what he or she already knows.

Although the framework focuses on how tools support individual learning processes differently, this support is not solely dependent on tool type. Instead, the learning mode enacted

and complexity of what is being learned affects the level of cognitive support desired in the learning process. For example, when individuals discover novel knowledge and, thus, radically change their mental models, high levels of support for intuiting and interpreting may be desired.

Finally, there are other theoretical lenses that may oppose the propositions derived from the theoretical framework developed here. For example, tool familiarity and expertise (Johnson and Marakas, 2000; Mackay and Elam, 1992), technological frames (Orlikowski and Gash, 1991; Orlikowski and Gash, 1994; Weick, 1990)⁷, adaptive structuration (Desanctis and Poole, 1994) and dominant design (Abernathy, 1978; Anderson and Tushman, 1990; Tushman and Anderson, 1986) theories may suggest different relationships and, thus, competing propositions. In the next section, I specifically address how technological frames provides a rival and competing theory to the one proposed above in terms of how tools affect knowledge discovery and learning.

3.3 Competing Theory

Using competing theories and the “logic of opposition” in the analysis of a phenomenon can lead to a richer and more complete understanding of that phenomenon (Robey and Boudreau, 1999; Webster and Watson, 2002). In this case, an understanding of how tools affect knowledge discovery and learning is sought. In contrast to the theory proposed above, a competing theory of technological frames is proposed. Technological frames was selected as a plausible competing theory based on its relationship to cognitive learning and mental models. This competing theory is discussed in detail below.

Technological frames are technology-oriented mental models, which involve the expectations, assumptions, norms and meanings associated with the use and outcomes of specific tools (Orlikowski and Gash, 1991; Orlikowski and Gash, 1994). Although these mental models change over time, they are sticky and resistant to change. The individual’s prior experience using specific tools, as well as the current dominant design of a particular tool class, influence the

⁷ While Orlikowski and Gash (1991; 1994) introduce and examine technological frames in detail, Weick (1990) also describes a form of technology-oriented mental model.

development of technological frames. A dominant design is the single design, architecture or product that has emerged as the product-class standard (Abernathy, 1978; Anderson and Tushman, 1990; Tushman and Anderson, 1986).

Technological frames affect future actions and use of a specific tool, as well as other new tools in the product class. Thus, technological frames affect and constrain both behaviours and cognition. When faced with an unfamiliar tool, individuals impose their technological frames of the “familiar” technology onto the “unfamiliar” technology (Hargadon and Douglas, 2001; Orlikowski and Gash, 1994). Despite any technical merits or specific functionality a tool may possess, tools most congruent with an individual’s technological frame will perform better than tools less congruent. This can be partially attributed to the individual’s familiarity with the tool that has most strongly contributed to the technological frame. In addition, when the individual attempts to use a new tool using their existing technological frame – akin to “unfaithful” appropriation (Desanctis and Poole, 1994) – high levels of incongruence will likely result in lower performance with the new tool.

Although dominant designs can only be known retrospectively (Anderson and Tushman, 1990), it can be argued that Google is the dominant design for the “Web-searching tools” product-class. This dominance is exemplified by people’s familiarity with Google and the fact that people use Google’s brand name as a verb – “I googled you”, “Google this term...”. Thus, when individuals think of searching on the Web, most have a sticky and similar set of expectations, assumptions and meanings associated with using search tools on the Web. Assuming Google is the dominant design, most individual’s technological frame regarding searching the Web will be strongly influenced by Google. As a result, tools that are most congruent with the Google technological frame will perform better than any other tool, despite technical merits or fit with the task. Thus, using the technological frames competing theory I propose:

P4a: Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all types of knowledge discovery (deep, broad and novel knowledge).

P4b: Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all modes of learning (mental model building, tuning and maintenance).

P4c: Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all individual learning process (information foraging, intuiting and interpreting).

By using a competing theory approach to analyze the forces that both facilitate and constrain how tools can be used for knowledge discovery and learning on the Web, this study helps to provide a richer understanding of external knowledge discovery and, more specifically, novel-knowledge discovery.

3.4 Summary

In order to address the two research questions identified at the beginning of this chapter, an external knowledge-discovery theoretical framework was developed in which a series of relationships between tool characteristics, type of knowledge discovered, and learning mode and processes supported, were proposed. A competing theories approach was adopted in order to better understand how opposing forces influence external knowledge discovery and learning on the Web. In addition, extensions to the 4I learning process model were proposed in the context of discovery and learning on the Web: 1) the addition of a fifth process – information foraging, and 2) the addition of a fourth level – the machine level.

Tools for the discovery of novel knowledge – one of the three levels proposed in the theoretical framework – are rare. To address this gap and research question 1, a novel-knowledge discovery (NKD) design theory is developed in Chapter 4.

CHAPTER 4

NOVEL-KNOWLEDGE DISCOVERY DESIGN THEORY

4.1 Introduction

As discussed in Chapter 3, tools to support the discovery of novel knowledge are rare. In order to address this gap and research question 1 – what design properties should a novel-knowledge discovery tool have in order to be effective? – a novel-knowledge discovery design theory is developed. In this chapter, the design-science approach to IT research is reviewed and the components of design theory discussed. Next, the novel-knowledge discovery (NKD) design theory is described in detail, including evaluation criteria and the theory-development process.

4.2 Design Science and Theory

Design science, also referred to as design theory, is an approach to research where artifacts designed to address an existing yet unsolved problem are built and evaluated (Hevner et al., 2004). Design theory is a prescriptive theory, based on social- and natural-science theories, that is specifically focused on goal achievement, in addition to explanation and prediction (Markus et al., 2002; Walls et al., 1992). Thus, design theory can be considered both a type IV theory – for explaining and predicting – and type V theory – for design and action (Gregor, 2006; Venable, 2006). Social-science theory and design theory are complementary theories. Social-science theory, the goal of which is to understand reality and develop theory, provides the theoretical rationale for the design principles used in design theory. Design theory, the goal of which is utility, adds relevance to social-science theory by creating useful artifacts for IS practice, which become the focus of social-science research (Hevner et al., 2004; March and Smith, 1995).

The utility of a design theory is defined as the effectiveness and usefulness of a design in addressing the stated problem and resulting requirements (Venable, 2006). Design theory is not a specific solution to one organization's business problem, but a solution to a class of problems

(Walls et al., 1992). Specifically, design theory addresses important unsolved problems in unique or innovative ways, or solved problems in more effective ways. Design theory provides a link between problem-space (a class of problems) and solution-space (a class of solutions). The theory created helps predict whether a design will solve the class of problems it was designed to address (Venable, 2006).

Design theory includes both a product of design and a process of design (Walls et al., 1992). The design theory deliverables are collectively referred to as the IT artifact. An IT artifact can be one of the following: *constructs*, the conceptualizations and vocabulary used to describe the problem or solution; *models*, to represent the problem space and provide abstractions of that space in order to envision solutions; *methods*, including processes or algorithms; and *instantiations*, or the actual implementation or prototype (Hevner et al., 2004; March and Smith, 1995). Thus, the definition of the IT artifact in design theory is fairly broad and includes the constructs, models, methods and instantiations “applied in the development and use of information systems” (Hevner et al., 2004 p. 82; March and Smith, 1995). In addition to technology-oriented IT artifacts, design theory can also focus on organizations, policies and work practices (Hevner et al., 2004; Venable, 2006).

A design theory, in addition to the creation of an IT artifact, consists of 1) *kernel theories*, from social- and natural-science theories, which illuminate the problem and drive design properties, 2) *meta-requirements*, the class of goals and problems to be addressed by the theory, 3) *meta-design*, a class of artifacts or set of design principles to address the meta-requirements, 4) *hypotheses* used to test the resulting theory, and 5) *design process guidelines* (Walls et al., 1992). Design principles, part of the meta-design, identify the “ideals to be achieved by the selection of specific technologies and the development of particular features [of the IT artifact]” (Boland et al., 1994 p. 464). Table 4.1 provides a sample of design theory exemplars, including a summary of the components included in the theory developed.

Table 4.1 – Design Theory Exemplars

Study	Kernel Theory	Meta-Requirement	Meta-Design	Design Process Guidelines	Evaluation
Bapna et al. (2004)	Auction theory	Improve design of economic information systems based on different use patterns.	<ol style="list-style-type: none"> 1) Developing user-centric bidding agents. 2) Inferring bidders' underlying valuations to facilitate real-time auction calibration. 3) Creating low-risk computational platforms for decision making. 	N/A	Supports design enhancements with evidence from data-driven taxonomy of bidding behaviour.
Boland et al. (1994)	Hermeneutics, inquiry systems	System to support distributed cognition.	<ol style="list-style-type: none"> 1) Ownership. 2) Easy travel. 3) Multiple perspectives. 4) Indeterminacy. 5) Emergence. 6) Mixed forms. 	N/A	Evaluated the system "Spider" against the design principles.
Boland and Tenkasi (1995)	Models of language, communication, cognition	System to support communication, knowledge work, facilitate perspective taking and perspective making.	<ol style="list-style-type: none"> 1) Creation of boundary objects (cause maps, narrative maps, models, classification schemes) to support perspective making and perspective taking. 2) Issue-specific space for perspective taking. 	N/A	Evaluated ideal systems against design principles.
Majchrzak et al. (2005)	Cognitive-affective model of communication	System to support collaboration know-how in distributed teams.	<p>Supports contextualization:</p> <ol style="list-style-type: none"> 1) Ownership. 2) Easy travel. 3) Multiple perspectives. 4) Indeterminacy. 5) Emergence. 	N/A	Evaluated existing systems against Boland et al.'s (1994) design principles.
Markus et al. (2002)	Organizational design as an emergent knowledge process	System to support emergent knowledge processes.	<ol style="list-style-type: none"> 1) Guidance for offline action. 2) Integrate expert knowledge with local knowledge sharing. 3) Componentized architecture. 	<ol style="list-style-type: none"> 1) Facilitate customer engagement. 2) Radical iteration with functional prototypes. 3) Dialectical development to reconcile rather than trade off conflicting requirements. 	Provided details about how the system (TOP Modeler) and the development process exemplify the 6 design principles. Evaluated the whole EKP process – both online use and offline use.

Study	Kernel Theory	Meta-Requirement	Meta-Design	Design Process Guidelines	Evaluation
Tillquist et al. (2002)	Resource dependency theory	Develop strategies for system development for complex organizational application domains.	1) Model the structure of dependencies in analysis and design to highlight problem areas. 2) Model political influence and economic resource exchange dependencies. 3) Modeling elements: activity, resource, role, goal, dependency, governance control.	N/A	Tool evaluated by applying it in a case study and comparing resulting model against design principles. Both "as is" and "to be" systems were modeled.

There is a lack of consensus on the elements comprising a design theory. For example, there is some debate regarding whether a design theory must contain both a design product and process. Walls et al. (1992) suggests the inclusion of both in a design theory. However, Venable (2006) separates the design product from the design process, suggesting that both represent discrete IT artifacts. The list of IT artifacts that a design theory may encompass, which includes constructs, models, methods, and instantiations, also differentiates the types of IT artifacts that can form the basis of a design theory (Hevner et al., 2004; March and Smith, 1995). In fact, the design process itself is a type of methods artifact. Thus, a design theory can focus on either the design product or a design process. In this study, I focus on the design principles surrounding the design product, not the design process. It is not clear that a new process for designing knowledge-discovery tools, used for novel-knowledge discovery, is needed. Instead, I am interested in the design principles for the class of knowledge-discovery tools used for discovering novel knowledge. As noted in Table 4.1, most design theories focus exclusively on the design product.

Two commonly used design-science frameworks include the Walls framework (Walls et al., 1992) and the Hevner framework (Hevner et al., 2004). These frameworks are complementary, yet focus on different aspects of design science. The Walls framework identifies the key elements of a design theory as noted above, whereas the Hevner framework describes the design-theory process and proposes a series of guidelines for evaluating design-theory research. The components, process and evaluation of the NKD design theory are discussed in detail below.

4.3 NKD Design Theory Components

The kernel theories related to the NKD design theory being developed in this study are the knowledge-discovery and learning theories discussed in the literature review chapter. These kernel theories provide the theoretical background to help explain the meta-requirements, or the challenges associated with discovering novel knowledge, and to derive the meta-design to address

these challenges. The following section describes the NKD design theory, including the meta-requirements and meta-design.

4.3.1 Meta-Requirements and Meta-Design

Meta-requirements describe the class of problem to be solved by a design theory (Walls et al., 1992). The meta-design describes the design principles that address the meta-requirements. In the case of the NKD design theory, these design principles address the challenges associated with novel-knowledge discovery. As described in Chapter 1, the problems associated with discovering novel knowledge are: 1) locating novel knowledge in the vast amount of information online, 2) recognizing novel knowledge as relevant, and 3) learning novel concepts. These problems, which form the meta-requirements for novel-knowledge discovery, and the associated design principles, are discussed in more detail below.

4.3.1.1 Problem 1: Difficulty Locating Novel Knowledge

The challenges associated with locating novel knowledge on the Web are influenced by two key factors: 1) the nature of the Web, and 2) the nature of novelty. The Web provides an important source of information for businesses (Beccerra-Fernandez et al., 2004; Chen et al., 2002; Chung et al., 2005). However, information overload is an ongoing problem for organizations that use the Web as a key source of information. Thus, locating novel knowledge is difficult because it is embedded within and interspersed throughout vast amounts of information on the Web, much of which is irrelevant to an individual or organization (Bowman et al., 1994; Chung et al., 2005). Text- and Web-mining techniques can help alleviate the problems associated with information overload (Chung et al., 2005; Rao, 2004) and potentially direct the organization's attention to novel knowledge. Using tools to more effectively capture an individual's attention for important information and knowledge has been noted as an important area for future knowledge management research (Davenport and Volpel, 2001).

Search engines are the most common method for discovering Web content. Search engines produce a long list of results, listing the Web sites that correspond to the search criteria, sorted by perceived relevance (Beccerra-Fernandez et al., 2004). The development of appropriate search terms is linked to successful location of information on the Web. Even when an individual knows what he or she is looking for, the initial search term(s) used may not provide the desired results, resulting in a refinement of search terms and an iterative search process. Thus, search-term precision is an important aspect of knowledge discovery on the Web. However, novel-knowledge discovery complicates the search process because, by definition of novel knowledge, the individual does not yet know what he or she is looking for.

Relevance criteria vary across search engines, but in general, are based on the similarity of the results to the query criteria (for example, keywords). The relevance criteria used in a search tool are critical because they influence the type of knowledge discovered and, thus, where the individual's attention is focused. The page-ranking algorithm, for example, uses a "global" computation to calculate the relevance of a page, based on how the rest of the world assesses that page's relevance (i.e., by linking to that page themselves). However, this is a weak approach because global relevance does not necessarily mean that the page is the most interesting and relevant in the context of the particular search. The long textual lists that are created by search engines using this type of relevance criteria can hide many interesting results and contribute to the problem of information overload described above (Chung et al., 2005). In contrast to search engines that use the page-ranking approach, search tools like Vivisimo¹ and Kartoo² organize and

¹ For example, Vivisimo (www.vivisimo.com) is a search and clustering tool that organizes and displays the results of a search in groups according to similarity of content, developing a taxonomy. Users can navigate and drill down into the taxonomy to find specific Web pages according to their interests.

² For example, Kartoo (www.kartoo.com) displays results in a knowledge map. A knowledge map represents the results in a graphical map, with nodes (circles) as Web sites and lines linking the nodes together to represent common key words between nodes. A knowledge map provides a visual summary of the results, allowing more details to be explored on demand (Chung et al., 2005).

display results in clusters and knowledge maps respectively in order to help the user find deeper structure amongst the Web pages returned.

The emergent and unpredictable nature of novelty suggests that it is challenging to search for a specific novelty. If one could search for a specific novelty, it would no longer be a novelty. As discussed in the theoretical framework in Chapter 3, the questions associated with novel knowledge are vague and unarticulable. Thus, the process of discovering novel knowledge also needs to be emergent, supporting scanning behaviours, and not reliant on relevance criteria and search-term precision.

4.3.1.2 Design Principles for Problem 1

I have identified three design principles to help address the challenges associated with locating novel knowledge. These principles are listed in Table 4.2 below.

Table 4.2 – Design Principles to Address Problem 1

Design Principles to Address Difficulty Locating Novel Knowledge	
Principle 1	Resolve information overload problem by distilling to a small result set.
Principle 2	Resolve information overload problem by directing the user's attention to the most significant result set.
Principle 3	Resolve information overload problem by reducing the reliance on search term precision through clustering.

One of the challenges inherent in information overload is the sheer volume of results returned to the user. Users have limited attentional resources and cannot process this amount of information effectively. In order to focus the user's attention, it is useful to limit the amount of "stimuli" presented to the user. One of the basic visual design guidelines suggests that providing a summary overview to users, and allowing them to filter and look at details on demand, improves usability (Shneiderman, 1996). Thus, distilling the result set returned to the user into a small, summarized overview is a useful approach for relieving the information overload problem.

However, providing a small, summarized list of results does not solve all aspects of the information-overload problem. As noted above, individuals have difficulty picking out the interesting and relevant results. Thus, the design of an NKD tool should also direct the user's

attention to the most interesting and significant results. Techniques such as clustering results by category and visualization may help the user explore the results and identify the most interesting results for their task (Shneiderman, 2000).

The next design principle addresses the issue of search term precision. As discussed above, locating information on the Web depends on specifying precise search terms, which is often challenging, especially when you do not know what you are searching for. Clustering, an approach where similar objects are grouped together, can help reduce reliance on precise terms. Rather than relying on the presence of an exact search term, clustering looks at similarity between documents and, therefore, documents that may differ slightly in terminology are grouped together. Thus, the design of an NKD tool should help reduce the reliance on search-term precision using clustering techniques. Reducing reliance on search-term precision through clustering, as well as the use of visualization, supports the exploratory and emergent processes associated with novel-knowledge discovery.

4.3.1.3 Problem 2: Difficulty Assessing Relevance of Novel Knowledge

Understanding and evaluating the relevance of search results is an important step in the knowledge-discovery process. One of the challenges with novel-knowledge discovery is that the individual, followed by the organization, must perceive the results of the discovery as relevant and important (Schulz, 2001; Vats and Skillicorn, 2004a). The uncertain relevance of new knowledge results in hard to predict and potentially remote implications. The novel knowledge discovered may not be relevant to one individual or area of the organization, but relevant to another (Amabile, 1983; Amabile, 1988; Csikszentmihalyi, 1996; Schulz, 2001; Witt, 2003). The results of a study examining organizational knowledge flows (Schulz, 2001) found that new knowledge flowed vertically within the organization so that it could be exposed to different knowledge and perspectives within the organization. This broad exposure not only highlights potential opportunities, but also risks and threats. Incremental knowledge, which has high

relevance, flowed horizontally to targeted areas within the organizations so that the knowledge could be exploited.

A study of knowledge reuse for radical innovation (Majchrzak et al., 2004), where discontinuous developments and significant improvements are made (Leifer et al., 2000), found that the most radical or novel solutions resulted when reusers found ideas in nontraditional sources that were outside their industry and dissimilar to their functional requirements and expectations. Although searching nontraditional sources is challenging in terms of understanding the relevance to the problem at hand, the result can be very relevant and beneficial. However, the challenge lies in understanding how the novel knowledge is relevant to the organization's business. For example, in the Reebok Pump™ case (Hargadon, 2002) described in Chapter 2, it may be challenging to see how inflatable splints and medical IV bags are relevant to the athletic shoe industry.

4.3.1.4 Design Principles for Problem 2

In order to address the problems associated with assessing the relevance of novel knowledge, two design principles for an NKD tool are proposed (see Table 4.3). These principles are discussed in more detail below.

Table 4.3 – Design Principles to Address Problem 2

Design Principles to Address Difficulty Assessing Relevance of Novel Knowledge	
Principle 4	Resolve uncertain relevance problem by providing metaknowledge for the result set.
Principle 5	Resolve uncertain relevance problem by providing measurement of "interestingness" to enable the evaluation of novel knowledge.

Providing access to metaknowledge, contextual clues surrounding the information, may help individuals assess the potential relevance of novel concepts (Majchrzak et al., 2004). In a previous study, metaknowledge surrounding a novel idea was used to evaluate the credibility, relevance and adaptability of the idea (Majchrzak et al., 2004). In fact, ideas were evaluated more positively by the mere presence of metaknowledge. Metaknowledge surrounding a concept could be used by individuals to understand and assess the relevance of the novel concept. Thus, the

fourth NKD design principle involves providing metaknowledge for the result set in order to facilitate an assessment of relevance.

In addition to metaknowledge, providing a measurement of interestingness for each result may support the evaluation of the novel knowledge and its potential relevance to the individual and organization. This measurement may provide the individual with additional clues regarding relevance and, consequently, whether to assess the information further with other group or organizational members. Thus, the fifth NKD principle involves providing a measurement of interestingness in order to enable the evaluation of novel knowledge.

4.3.1.5 Problem 3: Difficulty Learning Novel Knowledge

Learning novel concepts is difficult (Cohen and Levinthal, 1990) because the concepts are often outside of existing mental models, which are frameworks to help simplify and organize information (Crossan et al., 1999; Day, 2002; Hedberg, 1981; Kim, 1993). As discussed in the literature review, mental models guide interpretation and what the individual and the organization searches for and sees in the environment (Hargadon and Fanelli, 2002; Kim, 1993; Orlikowski and Gash, 1991). Novel knowledge is often outside existing mental models and therefore requires their modification. However, mental models are sticky and difficult to change. The stickiness of existing mental models constrains novel-knowledge discovery and learning novel concepts.

The challenges associated with learning novel knowledge are related to the properties of novelty, creativity and innovation. Innovation requires entrepreneurial intuition and imagination. Entrepreneurial intuition involves the ability to “make novel connections, perceive new or emergent relationships, and discern possibilities that have not been identified previously” (Crossan et al., 1999 p. 526). Similarly, imagination refers to the creation of something previously unthought of, rather than just “simple associations of adjacency” (Weick, 2006 p. 448). Both entrepreneurial intuition and imagination deal with exploration and discovering

something unknown and novel. The challenges associated with imagination and entrepreneurial insights in organizations stem from the difficulty in disengaging from what the individual and organization knows and the tendency to think incrementally rather than exponentially (Weick, 2006).

4.3.1.6 Design Principles for Problem 3

In order to address the challenges associated with learning novel concepts, two additional NKD design principles are proposed (see Table 4.4). A discussion of these design principles follows below.

Table 4.4 – Design Principles to Address Problem 3

Design Principles to Address Difficulty Learning Novel Concepts	
Principle 6	Resolve difficulty learning novel concepts by reframing existing knowledge.
Principle 7	Resolve difficulty learning novel concepts by enabling the examination of existing mental models by juxtaposing existing knowledge with indirectly related stimuli or knowledge.

As discussed in the literature review, there are competing views regarding the most effective way to change mental models and learn new concepts. These competing views include the following: 1) separating or discarding current mental models from the learning process (e.g. Hedberg, 1981; Kim, 1993; Nystrom and Starbuck, 1984; Tushman and O'Reilly III, 1999), 2) grounding the new knowledge in familiar knowledge and mental models (e.g. Hargadon and Douglas, 2001), and 3) reframing – conceptualizing existing mental models and knowledge in a new frame (e.g. Majchrzak et al., 2004; Watzlawick et al., 1974). The latter two views – grounding the new in terms of the familiar and reframing – leverage existing mental models in different ways. Grounding the new in terms of the familiar enables individuals to learn and make sense of new innovations, whereas reframing enables individuals to conceive of and make sense of novel ideas. Since this research is focused on the discovery and conception of novel ideas, reframing is a more appropriate approach in the context of novel-knowledge discovery. Grounding the new in terms of the familiar may be useful during the implementation phase – integration and institutionalization learning processes – of the innovation. Thus, the sixth NKD

design principle involves reframing existing knowledge in order to support learning novel concepts.

The discovery of novel knowledge and modification of mental models involves creating new mental concepts through a process of recombination. This process involves recombining existing cognitive concepts with external stimuli and new perceptions (Barnett, 1953; Hargadon, 2002; Witt, 2003). Juxtaposing existing knowledge with unrelated stimuli (Hender et al., 2002; Nagasundaram and Bostrom, 1995; VanGundy, 1988) or peripheral cases (Gogan, 2006) can help “shake up” the individual’s thinking and allow them to challenge their existing mental models (i.e. double loop learning). The result may be a novel and imaginative hypothesis or idea, as in the case of the Raynaud’s syndrome discovery. The challenge is to ensure the unrelated stimuli and peripheral cases are distant enough to identify imaginative possibilities, however not so remote that the individual is unable to see any connections and conceive of novel ideas (i.e. contextually appropriate). Thus, the seventh and final NKD design principle supports learning novel concepts by enabling the examination of existing mental models by juxtaposing existing knowledge with indirectly related stimuli or knowledge.

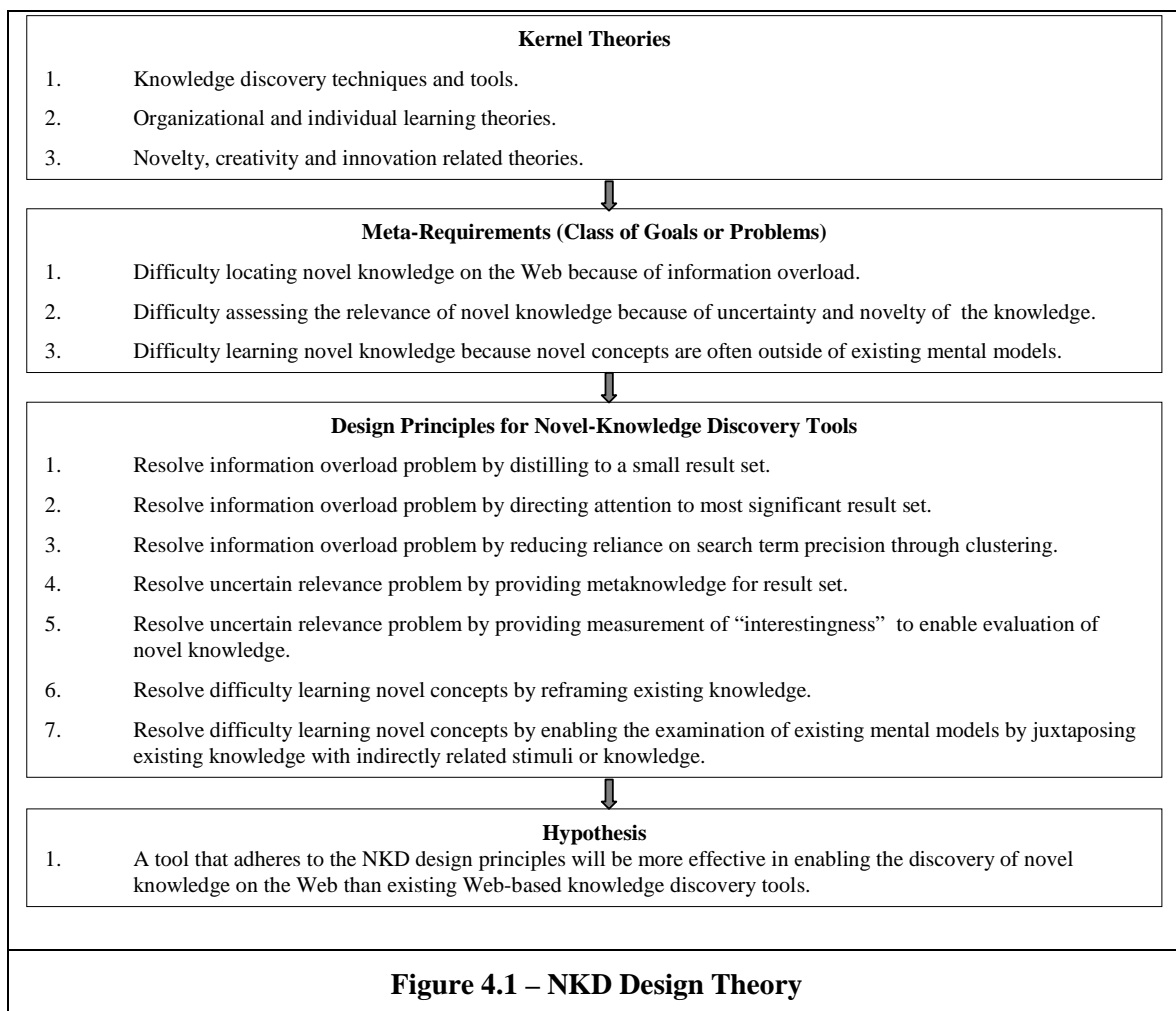
The following section discusses the evaluation of the NKD design theory.

4.4 Evaluation of a Design Theory

Design theory uses scientific theory as its kernel theories and uses the scientific method to test the resulting design theory (Walls et al., 1992). The feasibility of a particular design theory can be “supported by scientific theory to the extent that the design embodies the principles of the theory” (Walls et al., 1992 p. 38). Evaluation is built into the design theory itself through the development of hypotheses, which assess the degree to which the design principles address the meta-requirements. The effectiveness of a design theory is also commonly evaluated against the effectiveness of other technologies or designs (Venable, 2006). Accordingly, the proposition associated with the NKD design theory is as follows:

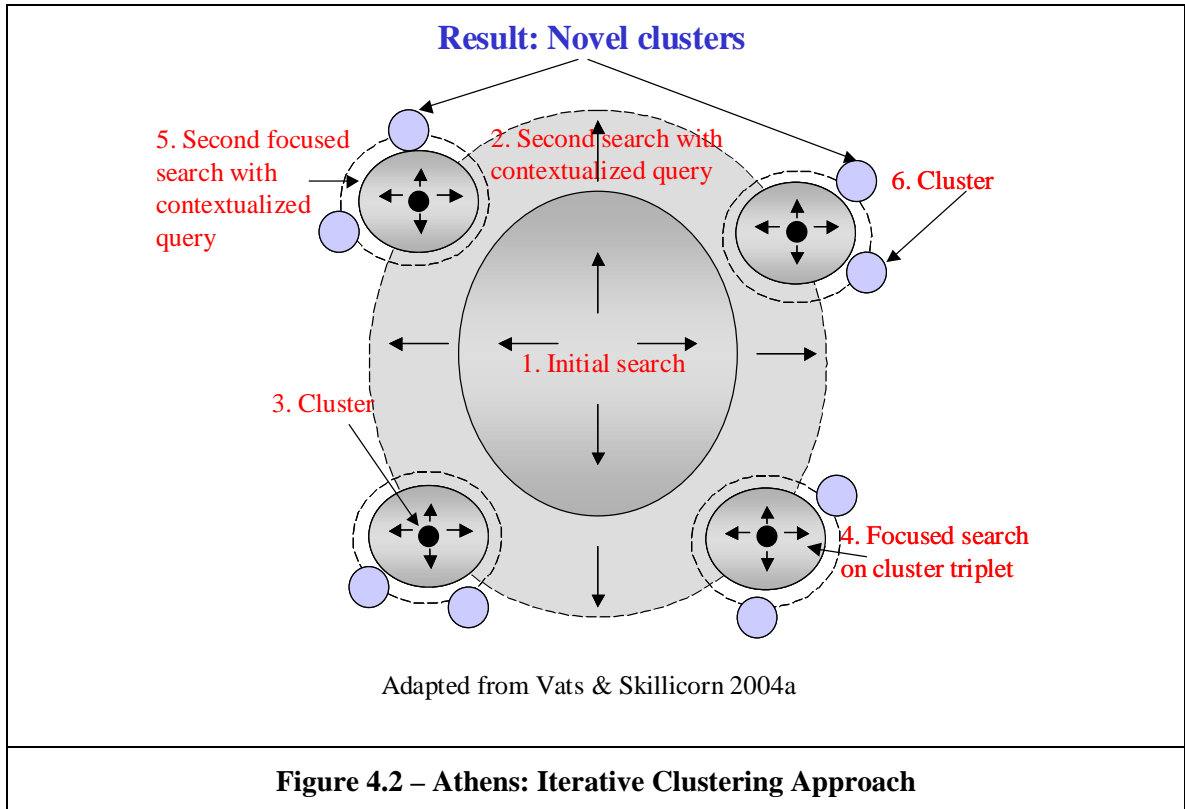
P5: A tool that adheres to the NKD design principles will be more effective in enabling the discovery of novel knowledge on the Web than existing Web-based knowledge-discovery tools.

Figure 4.1 provides a summary of the NKD design theory and the relationship between design components.



In a separate study, a tool – Athens 1.0 – was developed to discover novel knowledge prior to the development of the NKD design theory (Vats and Skillicorn, 2004a; Vats and Skillicorn, 2004b). The Athens tool uses an iterative clustering technique to discover novel knowledge on the Web, or other information repositories. Athens accepts several parameters, including keywords, Web site or domain specification (for example, IBM.com or *.uk), size of

results, and returns information that is indirectly related to the search terms, yet contextually appropriate. Figure 4.2 depicts the process employed by Athens to locate novel knowledge.



In the initial search step, Athens searches the Web based on the search terms entered, identifying concepts highly related to those search terms. Athens analyzes the retrieved pages, extracting the most important nouns related to the initial search terms. This step allows Athens to create background knowledge about the search terms – representing what you know about the topic – in order to further probe for novel knowledge. In the next step – second search with contextualized query – Athens performs another search of the Web, pairing the most important nouns extracted in the first search with the original search terms, to perform a highly contextualized query for knowledge connected to the original search terms. This step allows Athens to widen the scope of the search, while maintaining the context of the initial search terms. The third step – cluster – analyzes the results from this search and clusters them by similarity,

filtering out less-relevant results. The resulting “intermediate” clusters represent what you probably know about the topic. From these clusters, Athens extracts three nouns that best describe the cluster and uses these terms for another iteration of search, contextualized search, and clustering. This second iteration conducts a highly selective and focused query that results in a set of “novel” clusters that direct individuals and organizations to knowledge that is related to the search terms entered, but that is indirectly connected – representing what you probably do not know about the topic but may find interesting. Athens 1.0 labels and summarizes these clusters by a set of fifteen nouns that best describe the contents of the cluster.

As an illustrative example, had Don Swanson used Athens to discover novel knowledge about Raynaud’s syndrome rather than his manual approach, he would have entered search terms that represent what he knows – “Raynaud’s syndrome”. Athens would search for background knowledge about Raynaud’s syndrome, followed by a second contextualized search and clustering to determine what was probably known about Raynaud’s syndrome. One cluster would probably represent the subject of “blood viscosity” (intermediate level cluster). Athens would repeat this process a second time using these “intermediate” clusters as the starting point in order to find “novel” clusters – “fish oil” would probably be one of the novel clusters. Thus, Athens would create clusters that are two steps away from the initial topic of Raynaud’s syndrome. Similarly, in the MP3 case Athens would create novel clusters two steps away from the initial topic of MP3 technology. Gaming and the Wii product is one example of the possible contents of a novel cluster.

In the Athens 1.0 prototype, data-mining techniques are combined to create a new algorithm and method for novel-knowledge discovery. Data-mining principles and algorithms, such as clustering, singular value decomposition and spectral graph partitioning drove the initial design of this tool. In this dissertation, I use social-science theories as the kernel theories to develop an understanding of the meta-requirements for novel-knowledge discovery and develop the design principles that address these meta-requirements. Athens, the only tool that I am aware

of that at least partially addresses novel-knowledge discovery, was evaluated against the NKD design theory (see Table 4.5 below). The gaps identified in the initial prototype of Athens are addressed in a second version of the tool and evaluated experimentally (discussed in Chapter 6).

Table 4.5 – NKD Design Theory Principles and Athens

NKD Design Principles	Addressed?	Current Athens Version	Issues/Gaps
Resolve information overload problem by distilling to a small result set.	Yes	Allows the user to specify the number of results returned. Provides summary of novel clusters.	
Resolve information overload problem by directing attention to most significant result set.	Yes	Performs a selective and focused search. Clusters results in a set of distinct knowledge groups/areas.	
Resolve information overload problem by reducing reliance on search term precision through clustering.	Yes	Iterative search and clustering reduces reliance on specific terms. Mirrors and automates the manual process of searching and refining search terms multiple times.	
Resolve uncertain relevance problem by providing metaknowledge for result set.	Limited	Provides links associated with novel clusters, allowing users to better understand context of result set. Also, displays the three-word intermediate cluster, representing the path to the final novel cluster.	Onerous to click through all of the links for each novel cluster. Links are not integrated into the results page.
Resolve uncertain relevance problem by providing measurement of “interestingness” to enable evaluation of novel knowledge.	No	Not provided.	No interestingness measure provided. However, cluster label does sort 15 words by decreasing frequency.
Resolve difficulty learning novel concepts by reframing existing knowledge.	Limited	Underlying algorithms perform this reframing by searching down interesting paths related to the search term. However, difficult for the user to understand and interpret the output.	Challenging to understand the results and how they relate to search terms. Results are novel clusters with a 15-word label used to represent the cluster.
Resolve difficulty learning novel concepts by enabling the examination of existing mental models by juxtaposing existing knowledge with indirectly-related stimuli or knowledge.	Limited	Juxtaposes subject matter with indirectly related stimuli.	Results are currently difficult to understand, therefore difficult to examine existing mental models and learn novel concepts based on current form of output.

The criteria for a good design theory (Hevner et al., 2004) and how the proposed NKD design theory fulfills these criteria are summarized in Table 4.6 below.

Table 4.6 – Criteria Used to Assess Design Theory

Design Theory Criteria (Hevner et al. 2004)	This Study's Implementation
A viable artifact is produced.	Athens is an initial prototype produced prior to the development of the NKD design principles.
A relevant problem is addressed.	The discovery of novel knowledge can be strategically important to organizations, leading to learning capability advantages, innovation, or discovery of important strategic opportunities. Despite the significance of novel-knowledge discovery, organizations face a number of challenges in discovering novel knowledge. These challenges include: 1) difficulty locating novel knowledge, 2) difficulty recognizing novel knowledge as relevant, and 3) difficulty learning novel concepts that are outside of existing mental models.
The design artifact is evaluated.	Athens is evaluated against the NKD design principles. Gaps were identified and are being addressed in a second version of Athens. Further evaluation of the second version is completed in this research.
A research contribution in the area of design artifact, design foundations or design methodologies is made.	In this research, an NKD design theory is developed to better understand the path to discovering novel knowledge and overcoming the challenges of novel-knowledge discovery. Insights into how to learn novel concepts and change mental models will be gained.
Rigorous methods are applied in construction and evaluation.	Since the initial version of Athens was developed prior to the development of the NKD design theory, the initial evaluation of the design theory was against the current prototype. A number of gaps were identified, which are addressed in this research.
The design evolves, following a process such as the build-evaluate cycle.	The initial evaluation of Athens against the NKD design principles is the first step in an iterative build-evaluate design process. The gaps identified will be addressed and the tool will undergo further iterations of build and evaluate in this research.
The research is communicated to technicians and managers.	This research targets a managerial audience, avoiding extensive technical details regarding the instantiation of the tool. However, the method that Athens employs to discover novel knowledge is described.

Based on the NKD design principles identified above, the general characteristics of level 3 tools are added to the summary initially described in Chapter 3 (see Table 4.7).

Table 4.7 – General Characteristics of Knowledge-Discovery Tools

Tool Level	General Characteristics	Tool Examples
Level 1	<ol style="list-style-type: none"> 1. Searches across a domain of content using keywords entered by the user. (*, #) 2. Produces list of results matching search terms. (*) 3. Includes links to relevant content and meta-knowledge of that content. (*, #) 4. Sorts results by relevance / similarity to search terms (for example, weighted keyword scores, and page ranking). (*) 	Google, Yahoo, Altavista, Infoseek, Ask, Dogpile, Metacrawler ¹
Level 2	<ol style="list-style-type: none"> 1. Includes characteristics of level 1 tools (where noted *). 2. Groups results into clusters of similar content, thus identifying various dimensions of the search term and related concepts. (#) 3. Allows user to drill down into clusters for more detailed results. (#) 	Kartoo, Vivisimo, NorthernLight, Excalibur, Autonomy, Verity ²
Level 3	<ol style="list-style-type: none"> 1. Includes characteristics of level 1 and 2 tools (where noted #). 2. Distills and summarizes results to a limited set. 3. Directs user's attention to most interesting, relevant and novel results that are indirectly connected original search term. 4. Provides a measurement of interestingness to enable evaluation of results. 5. Reframes existing knowledge about the search term entered. 6. Supports users in challenging existing mental models regarding search term entered. 	Athens 1.0 (partially)
<p>Notes: ¹ Both Dogpile and Metacrawler have some features that are consistent with level 2, such as asking the user 'Are you looking for x?' and allowing the user to drill down into different search word combinations. ² Now part of Autonomy's product offerings.</p>		

The following section discusses the design-theory process.

4.5 Design-Theory Process

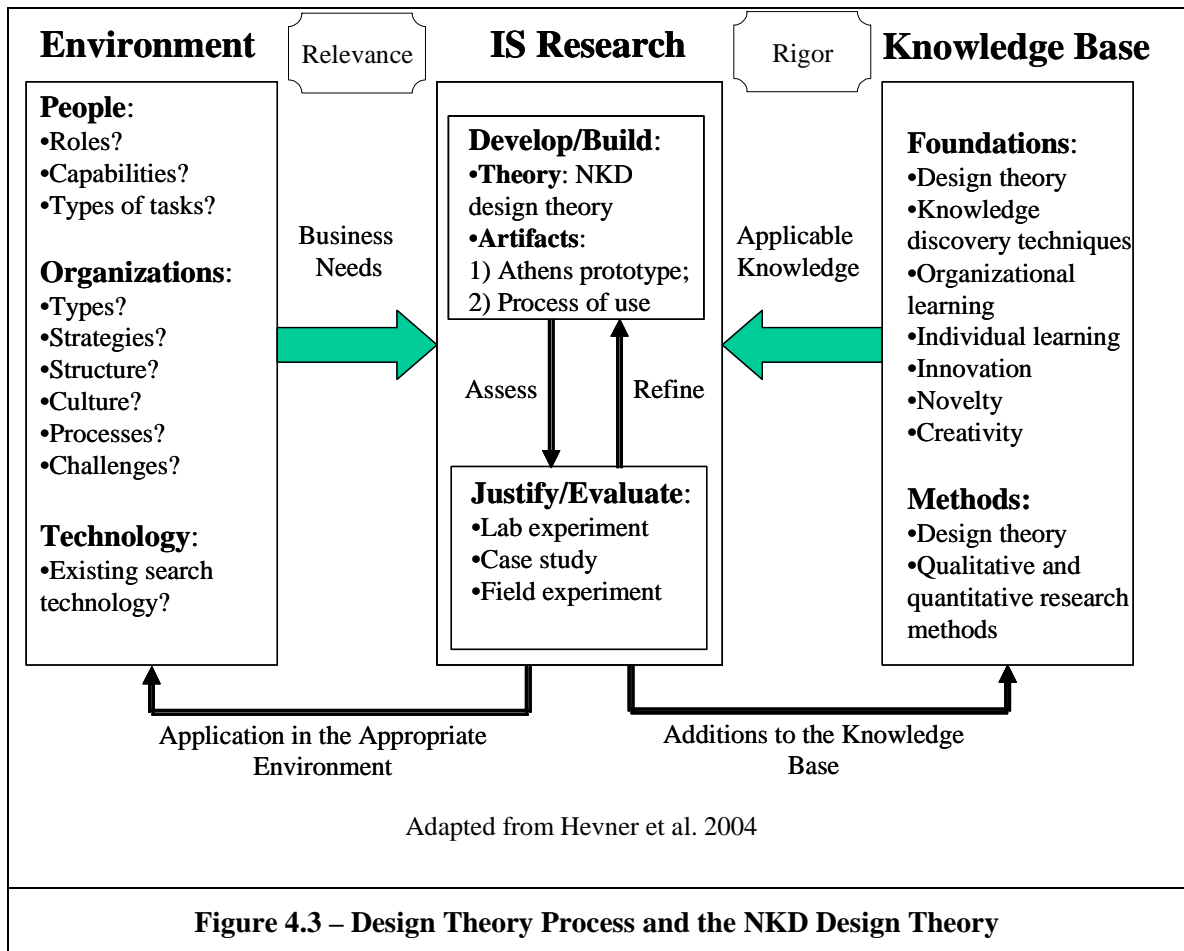
The design-science framework proposed by Hevner et al (2004) describes the design-theory development process as well as how behavioural science and design science interact. As depicted in Figure 4.3, there are three major components in the design-science framework: environment, IS research and knowledge base. The environment component defines the problem space – the business need and problem – including elements such as people, organization and technology. Design science is focused on developing a class of solutions to address the class of problem identified in the environment. Behavioural science also depends upon the environment to develop and justify theories about phenomena related to the problem space and business need. Thus, design science and behavioural science are driven from similar practical issues. The theoretical framework developed in Chapter 3 provides a partial understanding of the

environment and problem space in terms of types of tasks (people), individual and organizational learning processes and challenges (organization), and the types of technologies available (technology). Additional insights into the environment and the problem space can be gained through field studies, which are described in Chapter 6.

The knowledge base provides the foundational knowledge – existing behavioural and natural theories (kernel theories) – and methodological knowledge to inform the development of the design theory. The NKD design theory draws on a number of behavioural theories, as described in the literature review (see Figure 4.3 for a list of these theories).

For the third component in the framework – IS research – I use the knowledge of the problem space (environment) and the kernel theories from behavioural science (knowledge base), to develop the NKD design theory and artifact. Once the design theory is built, the theory and the artifact must be evaluated. The evaluation results in refinements to the design theory and an iterative build-evaluate cycle – a critical part of the design-theory process (March and Smith, 1995). In the case of the NKD design theory, I evaluated Athens – the prototype NKD tool – against the NKD design principles developed above. A number of gaps were identified, which were used to refine the Athens tool (Athens 2.0) and empirically evaluate the tool and theory using both qualitative and quantitative methods (described in Chapter 6).

Finally, design theory is meant to contribute to both the environment, by providing a useful artifact and knowledge for individuals and organizations, and the knowledge base, by updating either foundational or methodological knowledge. The design theory framework (Hevner et al., 2004) shows how design science can contribute both relevant and rigorous IS research.



4.6 Summary

To address the first research question posed in Chapter 3, an NKD design theory was developed in this chapter. An early prototype of the instantiation of the NKD design theory was described and evaluated against the design principles. Several gaps were identified, which are addressed in a new version of the tool and evaluated across the three studies included in this research – described in Chapter 6.

In addition to the instantiation of the NKD design theory described in this chapter, another design artifact is proposed in Chapter 5 – a process of using an NKD tool. Research question 2 – partially addressed in Chapter 3 – is further explored in Chapter 5 by examining the group processes involved in the use of an NKD tool.

CHAPTER 5

GROUP PROCESSES AND OUTCOMES

5.1 Introduction

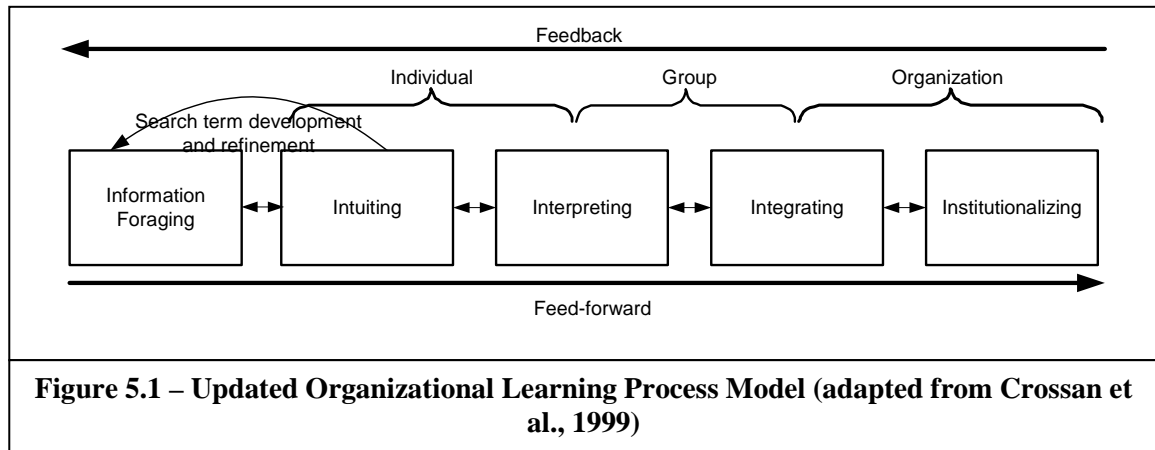
In the previous chapter, a set of design principles (meta-design) to address the challenges of novel-knowledge discovery were proposed. While these principles are important for developing a novel-knowledge discovery (NKD) tool, it is equally important to understand *how* such a tool should be used in order to be effective. Thus, the process of using an NKD tool to generate and evaluate ideas for the organization to pursue is examined and the second NKD design artifact created – process of use. This process includes both individual and group processes.

This chapter is organized as follows. First, an overview of group processes, support systems and outcomes is provided. Next, the outcomes evaluated in this research are described, followed by a discussion of three different processes of using an NKD tool to generate and choose novel ideas for the organization to pursue. A series of research propositions about these processes and their proposed outcomes are described.

5.2 Overview: Group Processes, Support Systems and Outcomes

In the theoretical framework developed in Chapter 3, I propose that NKD (level 3) tools support information-foraging, interpreting and part of the intuiting and search-term development processes, which have typically been individual-level processes. At the group level the critical learning processes include: 1) interpreting, where a shared understanding of the ideas is developed, and 2) integrating, where a decision is made regarding which idea to act upon and implement (see Figure 5.1 below). As with the individual processes, these group processes are challenging due to the novelty of the ideas generated and considered. The degree to which the group's decision and actions regarding the novel idea are integrated and later institutionalized within the organization depends on whether organizational members understand and agree upon

the new shared understanding of the novel idea. While the challenges regarding novel-knowledge discovery and learning by individuals are addressed by the proposed design principles, NKD tools do not directly support the group interpretation and integration processes that follow.



Processes at the group-level are supported by group support systems (GSS) (Desanctis and Gallupe, 1987; Sambamurthy and Poole, 1992). GSS tools “improve the process of group decision making by removing common communication barriers, providing techniques for structuring decision analysis, and systematically directing the pattern, timing, or content of discussion” (Desanctis and Gallupe, 1987 p. 589). Three levels of GSS tools have been identified, increasing in the level of structure and intervention into the natural group process. Level 1 tools provide structure for communication, such as anonymity and parallelism¹ (Dennis et al., 2001; Desanctis and Gallupe, 1987; Watson et al., 1988; Zigurs and Buckland, 1998); level 2 tools provide structure for information processing and consensus building, such as decision modeling (Dennis et al., 2001; Desanctis and Gallupe, 1987; Sambamurthy and Poole, 1992; Watson et al., 1988; Zigurs and Buckland, 1998); level 3 tools provide structure for information exchange within the group, such as controlling the pattern and timing of information exchange (Desanctis and Gallupe, 1987).

¹ For example, simultaneous input mechanisms.

In addition to the communication support and information-processing support structures provided by levels 1 and 2 tools, a third type of support structure has been identified – process structure (Dennis et al., 2001; Zigurs and Buckland, 1998). Process-structure support, also referred to as appropriation support (Dennis et al., 2001), involves defining the process the group will use to complete the group task, and ensuring that the group follows this process. GSS process structure features and mechanisms include agendas, facilitation, training and restrictive software (Dennis et al., 2001; Zigurs and Buckland, 1998). The focus here is on communication, information-processing, and process support for decision-making, which are all useful to the group learning processes of interpreting and integrating.

In order to explore how groups interpret novel ideas and decide which idea to act upon, I examine both the individual processes and group processes involved in generating and choosing ideas. Including both individual and group processes is congruent with other research in organizational learning (e.g. Crossan et al., 1999) and GSS (e.g. Potter and Balthazard, 2004).

Group research involving creativity tasks, such as brainstorming, and decision-making tasks has demonstrated that the outcomes from individuals performing these tasks and groups performing these tasks are different (Allport, 1920; McGrath, 1984; Potter and Balthazard, 2004; Steiner, 1972; Vinokur, 1969; Vinokur and Burnstein, 1974). Thus, I compare the outcomes of individual and group processes associated with using an NKD tool. How the NKD tool is used is proposed to influence the outcomes of tool use for the organization.

The outcome of these processes is a novel idea that has been selected for implementation by the organization – the integration process within the overall organization learning process. Outcomes of the creative process have typically been evaluated using criteria such as novelty and value or utility (Amabile, 1983; Couger and Dengate, 1992; Newell et al., 1962). Thus, the two outcomes under consideration in this research are: 1) the radicalness of the idea, and 2) the benefit of the idea to the organization. Following other GSS studies, the satisfaction and confidence with the decision made are also examined (e.g. Dennis et al., 2001; Fjermestad and

Hiltz, 1999; Gallupe and McKeen, 1990; Sambamurthy and Poole, 1992). Three processes of use are compared: 1) individual use of the NKD tool, 2) individual use of the NKD tool plus use of a GSS tool to filter and evaluate ideas, and 3) individual use of the NKD tool plus use of a collaborative-filtering tool to filter and evaluate ideas. These outcomes and processes are discussed in more detail below.

5.3 Outcomes

The outcomes of radicalness and benefit are discussed in the following section. A review of how these outcomes have been defined in past research, as well as a description of how they are defined and assessed in this research, are provided.

5.3.1 Radicalness

In general, radical is defined as “departing markedly from the usual or customary; extreme”². In the innovation and organizational literatures, radical ideas and innovations are described as controversial, and potentially disruptive (Hall and Martin, 2005); involving new processes or materials in the development of the idea (Barnett, 1953; Hill and Rothaermel, 2003); and, having technological content that is different from existing technology (Barnett, 1953; Dahlin and Behrens, 2005). The changes or improvements are discontinuous (Leifer et al., 2000), resulting in a product or service that differs substantially from the alternatives (Barnett, 1953; Zaltman et al., 1973).

Some authors have included benefit or impact in the definition of a radical innovation. For example, a radical innovation has been described as a new innovation that adds significant new value to the marketplace (Leifer et al., 2000). However, other authors separate technical content from impact (Dahlin and Behrens, 2005). The radicalness of the technical content and impact of the idea are two separate constructs that do not necessarily influence one another.

² From the American Heritage Dictionary, fourth edition.

Therefore, I measured the perceived radicalness of the idea separately from the perceived potential benefit of the idea.

5.3.2 Benefit

As discussed in the introduction, knowledge is defined as information with direction, enabling action-taking and decision-making and increasing the capacity for effective action. It is difficult to assess the benefit of knowledge without considering the opportunities for action and decision-making that knowledge affords. Thus, the benefit assessment in this study is focused on the ideas that are generated from the discovery of novel knowledge.

The potential benefit of an idea, opportunity or investment, for example, is often viewed as potential benefit less potential cost. Three types of benefits and costs have been identified in the literature: 1) tangible – costs and benefits that are directly measurable, for example increased variable costs; 2) quasi-tangible – costs and benefits that have some directly measurable aspects and other aspects that are difficult to measure and quantify; for example, improved shipment accuracy may lower costs and improve customer satisfaction; and 3) intangible – costs and benefits that are indirectly measurable and thus require a subjective assessment, for example improved morale (Parker et al., 1988; Ryan and Harrison, 2000). Table 5.1 provides a description of the factors that organizations commonly consider when evaluating opportunities (Harrison and St. John, 2004).

Table 5.1 – Decision Factors in Opportunity Evaluation (from Harrison and St. John, 2004)

Factor	Description
Cost/benefits analysis	The degree to which the financial benefits of pursuing the opportunity outweigh the financial costs.
Ethical analysis	Whether the opportunity under consideration is consistent with the organization's strategy and whether the organization's reputation could be negatively affected.
Protection of organizational strengths	How the pursuit of the opportunity will affect the organization's other strengths.
Implementation ability	The degree to which the implementation of the opportunity is congruent with the capabilities, structure, systems, processes, and culture of the organization.
Stakeholder analysis	How the opportunity will affect key stakeholders, whether they will approve the decision, how they will react to the implementation of the opportunity and the ramifications of this reaction.
Future position	The viability of the opportunity as the industry and environment evolves. Evaluates whether the opportunity is required for survival or competitive success.

Each of these factors can be interpreted as a cost and/or benefit of the idea under consideration and range from tangible to intangible in nature. The importance of each factor will vary across opportunities based on the types of activities involved. For the purposes of this research, the potential benefits of the novel idea for the organization are assessed in light of costs, including those that are tangible, quasi-tangible and intangible.

5.4 Processes

The process for novel-knowledge discovery and associated organizational learning is proposed to begin with an organizational problem or question. Group members begin the individual learning process by forming their own interpretation of the problem and formulating one or more keywords to be used in the information-foraging process. The NKD tool supports the information-foraging, interpretation and, partially, the intuiting processes, including further refinement of the search terms. Individuals learn and reframe their existing mental models based on the tool output. This output serves as “unrelated stimulus”, which results in each individual generating highly novel ideas. The novel ideas generated by each individual, divergent at first, are combined, ranked, reviewed and discussed by the group. The group attempts to come to a shared understanding regarding these ideas and choose the best one to pursue.

The challenges faced by individuals in discovering and understanding novel knowledge were discussed in Chapter 4. In addition to these individual challenges, there are a number of challenges with novel-knowledge discovery at the group level, which are discussed in detail below.

5.4.1 Novel Knowledge and Group Challenges

There are a number of challenges associated with developing a shared understanding of novel knowledge within groups composed of individuals with a broad range of organizational perspectives and knowledge. However, this broad range of perspectives and knowledge is important to fully assess the potential relevancy of the novel ideas for the organization (Schulz,

2001). The group may have interpretive differences regarding how they ascribe meaning to the novel ideas. Interpretation is constrained by an individual's prior knowledge and mental models. Thus, the group must attempt to overcome these differences in interpretations in order to come to a shared understanding (Carlile, 2004). The process of coming to a shared understanding requires a certain amount of common knowledge between group members in order to communicate effectively and leverage the specialized knowledge of each group member (Carlile, 2002; Carlile, 2004; Cohen and Levinthal, 1990; Grant, 1996). However, when dealing with novel ideas, the common knowledge across the group may no longer be effective to understand and communicate how specialized knowledge applies to the novel idea. Thus, while one member of the group may understand a particular novel idea and how it could be valuable to the organization, other members of the group may have difficulty sharing this understanding of the novel idea.

Another challenge for the group in coming to a shared understanding and, more so, for reaching a consensus regarding an idea to pursue, is the knowledge that each individual has "at stake". Each member of the group has his or her own set of knowledge and expertise acquired, and own set of interests. Members view this knowledge and their interests as "at stake" due to the costs of discarding existing knowledge (Hargadon and Fanelli, 2002). Novel ideas may be particularly costly to certain group members whose knowledge is "at stake" (Carlile, 2004).

The increased social complexity – due to the coordination required – for group tasks affects the creativity and imagination of the group. In order to coordinate, group members tend to rely on their common knowledge and existing relationships within their shared mental models rather than imagining novel possibilities (Weick, 2006). As a result, highly novel ideas may be discarded in favour of less radical ones. In addition, any new ideas created by the group as a result of group brainstorming synergies will tend to be no more novel than individual ideas.

The influence that increased social complexity has on the radicalness of the group's ideas may be moderated by the presence of a minority. Prior research has shown that when a minority

viewpoint is expressed in groups, divergent thinking, number of alternatives considered and the novelty of the ideas increases (Nemeth, 1986; Ocker et al., 1996).

In the section below, I compare the outcomes of three different individual and group processes associated with using an NKD tool to generate and choose an idea for the organization to pursue. These processes include: 1) individual use only, 2) individual use and GSS tool, and 3) individual use and collaborative-filtering tool.

5.4.2 Individual Use Only

In the first process – individual use of the NKD tool – an individual uses the NKD tool to search for and generate novel ideas addressing the specified organizational task. From their list of ideas, the individual selects the best one for the organization to pursue. There are two opposing forces that affect the number and creativity of the ideas generated by individuals and groups. Group synergies will tend to result in both higher quantities and levels of creativity of ideas generated by groups than individuals (e.g. Allport, 1920; Osborn, 1953). However, process losses counteract this effect. For example, the creativity literature suggests that individuals generate more, as well as more creative, ideas than groups due to process losses such as production blocking (McGrath, 1984; Potter and Balthazard, 2004; Steiner, 1972) and social complexity (Weick, 2006).

The relatedness of the stimuli affects the cognitive complexity of this task and thus the number of ideas generated (Hender et al., 2002; VanGundy, 1988). The NKD tool provides individuals with “unrelated stimuli” as a result of the indirectly connected and novel clusters produced. As a result, each individual will produce a relatively small number of ideas as opposed to the large numbers typical in brainstorming activities (Osborn, 1953).

5.4.3 Individual Use and GSS Tool

In the next process under consideration – individual use of the NKD tool plus the use of a GSS tool – an individual uses the NKD tool to search for and generate novel ideas addressing the

organizational task. The ideas are combined into one list and provided to the group to evaluate individually, without interaction. The individual scores are then combined and the ideas – content and scores – are discussed with the group members to build consensus. After the discussion, a final vote is taken to help the group come to a consensus regarding the “best” idea to pursue. Thus, this process uses two rounds of group evaluation, mediated by group discussion.

In the group literature, the activities included in this process include the generation of ideas, followed by a decision-making task where there is no “correct” answer³ and consensus on a preferred alternative is required (McGrath, 1984). Past research has examined the fit between the different types of group tasks and the GSS structures provided (e.g. Dennis et al., 2001; Zigurs and Buckland, 1998). The fit between the task and the GSS structures utilized for the task has been shown to increase the effectiveness of the group, whereas the provision of process structure increases the efficiency and satisfaction with the process (Dennis et al., 2001). Group tasks that include both the generation of ideas and decision making with consensus achieve better fit when electronic communication is used for the generation of ideas and when electronic and verbal communication plus information-processing support are used for the decision-making task (Dennis et al., 2001). Further, GSS are designed to minimize process losses such as production blocking and social complexity so that group synergies are maximized.

In this case, the GSS system provides electronic communication support for idea creation and decision making – ideas displayed, voting anonymously – and information-processing support for decision making – evaluation results are displayed in a structured format, such as a decision model. In addition, verbal communication is included in the decision-making process to enable the group to develop a shared interpretation surrounding the ideas and resolve any conflicting points of view. Thus, this process should help address some of the challenges associated with groups and novel ideas discussed above.

³ For example, in mathematical models, there may be a “correct” solution.

Research on group decision-making tasks has found that groups tend to shift their original preferences after the group discussion towards either a more risky or more conservative alternative (e.g. Stoner, 1968; Vinokur, 1969; Vinokur and Burnstein, 1974)⁴. In the context of this study, risky can be thought of as a more radical idea and conservative as a less radical idea. The direction of the group's choice shift depends on the degree to which partially shared persuasive arguments are made during the group discussion, where the group majority lies and ambiguity in the alternatives (Vinokur, 1969; Vinokur and Burnstein, 1974). During the group discussion, if partially shared persuasive arguments are made regarding a risky or conservative idea, group members who initially preferred an idea towards the opposite end of the risky-conservative continuum will tend to shift their opinion towards this other idea based on the new persuasive arguments. Group majority has no impact in this case. Alternatively, if persuasive arguments do not exist for any of the ideas – as a result of ambiguity in the ideas – then individuals will tend to conform to the group majority, based on social comparison theory (McGrath, 1984; Vinokur, 1969). Thus, groups have a tendency to make either riskier or more conservative decisions than individuals (McGrath, 1984; Vinokur, 1969; Vinokur and Burnstein, 1974)⁵.

5.4.4 Individual Use and Collaborative-Filtering Tool

In the last process under consideration – individual use of the NKD tool plus the use of a collaborative-filtering tool – an individual uses the NKD tool to search for and generate novel ideas addressing the organizational task. The ideas are combined into one list and provided to the group to evaluate individually, without interaction. Individual scores are then provided to a collaborative-filtering tool to “automatically” choose the idea to pursue.

⁴ Stoner's unpublished Master's thesis from 1961 is widely cited as the initial source of this finding.

⁵ Note - early group research suggested that groups tend to make riskier decisions, known as the risky-shift phenomenon (Vinokur, 1969).

A collaborative filtering tool is a data-mining technique that combines the opinions or judgments of multiple individuals, evaluates and weights each individual's ability to judge, and provides a "global" recommendation (Skillicorn, 2001; Skillicorn, 2007). Google uses a similar approach in page-ranking, where a Web page's relevance is assessed by the number of pages that link to it, weighted by the quality of those pages (Chen, 2001). Collaborative filtering applies ensemble learning principles, where collections of individual classifiers that are diverse and accurate are constructed (Dietterich, 2000). Each individual classifier in the ensemble votes for the decision they deem best, resulting in a highly accurate global classification decision (Dietterich, 2000). The collaborative-filtering tool in this study uses singular-value decomposition⁶, scripted in Matlab.

In contrast to a democratic voting system, which provides equal weight to each member's perspective, collaborative filtering discounts the opinions of individuals who appear to be "weak" assessors and emphasizes the opinions of individuals who appear to be "strong" assessors. The quality of an individual assessment is based on the degree to which an individual's opinions are in agreement with the opinions of other individuals with better judgment skills (Skillicorn, 2001). Collaborative filtering resolves this circular relationship to provide a collective recommendation that retains the variance in opinions across the group and develops a bias towards strong assessors. One can think of a strong assessor as a "discerning" evaluator. Movie ratings provide a useful example. On a continuum of best to worst movies, the best movies are known as "blockbusters". These movies are universally popular – most people like the movie. The worst movies are also universally viewed as bad – otherwise known as "B" movies. However, it is often more interesting to look at movies that are liked by a relatively smaller number of discerning movie critics. These movies don't appeal to everyone, but are viewed as great movies by a small number of individuals with high-quality movie-rating skills. Thus, the collaborative-

⁶ Singular-value decomposition is an unsupervised data-mining approach that rotates a multi-dimensional dataset (matrix) so that variance is maximized on the first axis and remaining variance is subsequently maximized on the other axes (Skillicorn, 2007).

filtering tool distinguishes between these two types of raters: 1) raters who show an ability to discern different “grades” of movies, similar to other “high quality” raters, and 2) raters who like or dislike most things. The “high quality” raters are assigned more weight so that the global recommendation is skewed towards the movies, or ideas in this case, that these raters evaluate highly. Thus in theory, the results from a collaborative-filtering tool should differ from a democratic vote.

A comparison of the proposed outcomes of each of the three processes discussed above is provided below.

5.4.5 Comparison of Use Processes

Three use processes were described above: 1) individual use, 2) individual use and GSS tool, and 3) individual use and collaborative-filtering tool. The outcomes associated with individual use of the NKD tool are limited by the individual’s interpretation of both the problem and relevancy of the results. As suggested by Schulz (2001), the relevancy of new knowledge is best assessed by multiple individuals with diverse knowledge. In addition, the number of ideas generated and, thus, considered in the “individual use only” process is less than the other two processes where the ideas of all individuals are combined. Thus, the range of ideas under consideration is narrower than the other two processes. Compared to the second and third processes, which include the sharing of multiple perspectives, the first process – NKD tool use by individuals only – will probably not produce ideas with as much radicalness or organizational benefit, on average. In other words, the probability of one particular individual producing an idea with as much radicalness and benefit as a group of individuals is low.

In the “individual use and GSS tool” process, the ability to discuss ideas from multiple perspectives and address any partially shared persuasive arguments should result in the selection of an idea that is either more radical or more beneficial than those made in the other two processes. The partially shared perspectives will, thus, result in a choice shift from the group’s

initial position. A more radical idea may be selected if persuasive arguments about a particularly radical idea are made that shift other group member's opinions during the discussion. It's unlikely that the group will chose an idea that has little benefit. Thus, the radical idea chosen will be beneficial as well. In contrast, a more beneficial idea may be selected if persuasive arguments about a highly beneficial, less radical idea are made that shift other group member's opinions during the discussion

The main difference between the "individual use and GSS tool" process and the "individual use and collaborative-filtering tool" process is the group discussion resulting in consensus⁷. In the GSS process, consensus-building through debate and discussion reduces variation in the group through persuasive arguments, creating a high level of agreement. The subsequent democratic and anonymous vote "eliminates" any biases. In the collaborative-filtering process, it is unlikely there will be a high level of agreement. However, the strong assessors will tend to be in agreement regarding the "good ideas" (high on both radicalness and benefit) and the "bad ideas" (low on benefit). The bias towards the strong assessors will result in an idea being selected that is high on both radicalness and benefit. While the persuasive arguments made during the group discussion in the GSS process may result in the same outcome as the collaborative-filtering process, it is argued that the persuasive arguments made in the group discussion will likely favour more beneficial ideas. For example – although social complexity and coordination issues are addressed by GSS tools through features such as voting, the group discussion and debate may result in coordination issues that move the group away from controversial and radical ideas (Weick, 2006). Thus, the propositions are as follows:

P6: The process used with level 3 knowledge-discovery tools will influence both the benefit and radicalness of the idea chosen.

⁷ One could argue that information-processing support also differentiates these two processes. However, one could also argue that collaborative filtering is a form of information-processing support given its decision modeling properties.

P6a: Level 3 knowledge-discovery tools combined with individual and collaborative-filtering processes will choose ideas that are more “radical” than individual and group processes alone.

P6b: Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building, will choose ideas that are more “beneficial” than individual or collaborative-filtering processes alone.

These propositions imply that the process of using an NKD tool has tradeoffs in terms of outcomes for the organizational learning process: radicalness of the idea versus potential benefit.

Lastly, in terms of satisfaction with the process and confidence in the decision made, it is proposed that the GSS process, which involves discussing the ideas, reaching a consensus based on group discussion, and higher levels of process support⁸ will result in higher satisfaction and confidence scores. Thus,

P7a: Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building and process support, will result in higher levels of satisfaction with the process than individual or collaborative-filtering processes alone.

P7b: Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building and process support, will result in higher levels of confidence with the decision made than individual or collaborative-filtering processes alone.

5.5 Summary

In this chapter, the group learning processes associated with novel-knowledge discovery were examined. Specifically, three different processes of using an NKD (level 3) tool to generate and choose an idea for the organization to pursue were described. These processes combine individual and group processes, and include the use of GSS and collaborative-filtering tools. Research propositions regarding the outcomes – radicalness, benefit, satisfaction and confidence – of these processes were developed. This chapter helps address research question 2 and develop a second design artifact for the NKD design theory – process of use. Chapter 6 describes the research design that is used to evaluate the research propositions identified in Chapters 3 through 5.

⁸ Due to the more active role played by the facilitator in process 3.

CHAPTER 6 RESEARCH DESIGN

6.1 Introduction

In order to evaluate the propositions identified in Chapters 3 through 5 (see Table 6.1 for a summary), the research design was composed of three studies. Prior to the beginning of these studies, the NKD tool – Athens – was modified to align with the NKD design theory, addressing the gaps identified in Chapter 4. In Study 1, a lab experiment was conducted to test the NKD design theory as well as the propositions that differentiate knowledge-discovery levels and tools and how they support learning. In Study 2, a revelatory single-setting case study was conducted to explore novel-knowledge discovery in practice, including an evaluation of the NKD tool Athens. In the third study, a field experiment was conducted to evaluate the processes of using an NKD tool, with a specific focus on group learning processes.

In the sections below, the experiments (Studies 1 and 3) and the case study (Study 2) are described in detail, including evaluation criteria.

Table 6.1 – Summary of Propositions

#	Proposition
P1	Knowledge-discovery levels can be differentiated by the type of knowledge sought.
P2	Associated levels of knowledge-discovery tools exist.
P2a	Level 1 knowledge-discovery tools ¹ most effectively discover incremental details about what is already known, in order to deepen one's understanding of a concept.
P2b	Level 2 knowledge-discovery tools ² most effectively discover multiple dimensions of a concept and relationships with other concepts, in order to broaden one's understanding of a concept.
P2c	Level 3 knowledge-discovery tools ³ most effectively discover novel knowledge and surprising associations.
P2d	Level 1 knowledge-discovery tools support mental model maintenance.
P2e	Level 2 knowledge-discovery tools support mental model tuning.

¹ Such as Google.

² Such as Vivisimo.

³ Such as Athens.

#	Proposition
P2f	Level 3 knowledge-discovery tools support mental model building.
P3a	Level 1 knowledge-discovery tools ⁴ are able to support the information foraging process.
P3b	Level 2 knowledge-discovery tools ⁵ are able to support the information foraging process and partially support the interpretation process.
P3c	Level 3 knowledge-discovery tools ⁶ are able to support the information foraging and interpretation processes, and partially support the intuiting process, including search-term development.
P4a	Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all types of knowledge discovery (deep, broad and novel knowledge).
P4b	Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all modes of learning (mental model building, tuning and maintenance).
P4c	Tools that are most congruent with the Google technological frame will perform better than any other type of tool, across all individual learning process (information foraging, intuiting and interpreting).
P5	A tool that adheres to the NKD design principles will be more effective in enabling the discovery of novel knowledge on the Web than existing Web-based knowledge-discovery tools.
P6	The process used with level 3 knowledge-discovery tools will influence both the benefit and radicalness of the idea chosen.
P6a	Level 3 knowledge-discovery tools combined with individual and collaborative-filtering processes will choose ideas that are more “radical” than individual and group processes alone.
P6b	Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building, will choose ideas that are more “beneficial” than individual or collaborative-filtering processes alone.
P7a	Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building and process support, will result in higher levels of satisfaction with the process than individual or collaborative-filtering processes alone.
P7b	Level 3 knowledge-discovery tools combined with individual and group processes, such as consensus-building and process support, will result in higher levels of confidence with the decision made than individual or collaborative-filtering processes alone.

6.2 Experiments

6.2.1 Overview

In general, the purpose of an experiment is to test predictions derived from theory and to study relationships under controlled conditions where confounding conditions can be minimized.

There are two major types of experiments: laboratory and field (Benbasat, 1990). The defining

⁴ Such as Google.

⁵ Such as Vivisimo.

⁶ Such as Athens.

characteristics of the laboratory experiment are as follows: 1) setting is contrived and artificial, 2) participants are assigned randomly to treatment and control conditions, 3) one or more independent variables are manipulated and the impact on the dependent variable of interest is measured, and 4) all other independent variables and intervening variables are controlled so that causal inferences can be made (Stone, 1978). Benbasat suggests that MIS researchers can leverage the artificial setting of laboratory experiments by using them to develop and test novel systems that do not have a real life counterpart. Field experiments are similar to laboratory experiments except: 1) the setting is in the field and thus natural instead of artificial, 2) there is less control over the other independent variables and intervening variables, and 3) participants are usually not assigned randomly, resulting in a quasi-experiment (Campbell and Stanley, 1963; Shadish et al., 2003). As a result of the confounding variables and non-random assignment, causality is harder to establish and there are more rival explanations for the results that need to be ruled out. However, field experiments benefit from the natural setting, increasing the generalizability of the results.

The main benefit of experiments is that they are very strong in enhancing internal validity, which refers to the validity of inferences made regarding causality – a change in the independent variable causes the change in the dependent variable. Internal validity depends on the extent to which one can rule out rival explanations for the observed change in the dependent variable. These “threats” to internal validity can and should be reviewed prior to experimental design and the appropriate design features included to eliminate and mitigate these threats (Benbasat, 1990).

A major weakness of experiments is external validity, which refers to the generalizability of the findings to other settings (Benbasat, 1990; Shadish et al., 2003). One of the tradeoffs in experimental design is between internal and external validity, or control and realism (Mason, 1990). Internal validity is gained through the control of the variables and setting; external validity is gained through the realism of the setting.

6.2.2 Key Design Elements of Experiments

There are several important elements to consider when designing experiments, including: the method of assignment, as well as the use of pretests, posttests, and control groups. These design elements are discussed in more detail below.

A key characteristic and differentiator of experiments is the method of assigning participants to conditions. In randomized experiments, participants are assigned to treatment or control conditions using a random process. In other words, assignment to a condition is based on chance. The likelihood of being assigned to a particular condition is equal across all participants. As a result, the groups of participants assigned to each condition are “probabilistically similar to each other on the average” (Shadish et al., 2003 p. 13). As a result, any differences between these groups on the dependent variable measured at the end of the experiment are more likely to be the result of the different treatment conditions, not inherent differences in the group. Random assignment eliminates selection bias – a threat to internal validity that arises when participants assigned to one treatment condition differ from participants assigned to another treatment condition and these differences, for example IQ, relate to the variance in outcomes between groups.

In some cases, random assignment to treatment conditions is not possible or ethical. Experiments that do not randomly assign participants to conditions are known as quasi-experiments (Campbell and Stanley, 1963; Shadish et al., 2003). In quasi-experiments, participants can be assigned to treatment conditions based on self-selection or the decision of the researcher. Selection bias is a plausible threat to internal validity in the case of quasi-experiments; therefore, the researcher must take this rival threat into account.

A pretest that measures the same outcome variables as the posttest allows for a comparison of initial group differences prior to the introduction of the treatment conditions. A pretest is therefore useful when random assignment is not possible to rule out selection bias. In addition, comparing the pretest scores to posttest scores allows the researcher to assess the

magnitude of the treatment effect on the dependent or outcome variables. In some research designs, multiple pretests are conducted to help rule out maturation as a threat to internal validity (Shadish et al., 2003). Pretests may not always be feasible and in some cases, may sensitize participants to the treatment and affect the results of the study (Lana, 1969; Shadish et al., 2003).

Posttests involve measuring the dependent variable after the treatment intervention in order to measure the effect of the treatment. In some research designs, multiple posttests are conducted in order to assess the permanency of the effect or to measure a delayed effect, if applicable to the study (Shadish et al., 2003). Posttests are critical because they help identify whether the treatment had an effect or not.

A control group is a comparison group that is compared to the treatment group and either receives no treatment or an alternative treatment. Multiple control or comparison groups may be used in a design (Shadish et al., 2003). Posttest results of all groups – treatment and control – are then compared to measure and compare the effects of each type of treatment.

In any research design, an important goal is to control several different types of error variance: 1) variance in the independent variable, 2) variance attributed to extraneous or confounding variables that may affect the results, and 3) variance attributed to measurement error (Stone, 1978). In experiments, the manipulation being tested should maximize the variability of the independent variable across treatment conditions in order to maximize the effect on the dependent variance. If there is insufficient variance in the independent variable, an effect may not be detected or not appear significant.

Extraneous or confounding variables are those that covary with the outcome or dependent variables. If these variables vary across treatment conditions, they represent an alternative and rival explanation to the hypothesized relationship being tested. Variance in confounding variables should be controlled using one of the following methods: 1) randomly assign participants to treatment conditions so that variance in confounding variables is equal across treatment groups, 2) measure confounding variables in order to control for their effects, 3) hold

the confounding variable constant – if there is no variance in the confounding variable, it cannot be linked to effects of the independent variable on the dependent variable (Stone, 1978).

Lastly, variance in the dependent variable attributable to measurement error should be minimized. Two sources of measurement error include: unreliable measures and noise. Increasing and measuring the reliability of scales used in the study is one way to reduce error variance. This can be accomplished by creating a composite of items to represent each measure, where feasible, and measuring the composite's internal consistency. The more items in a composite, the more reliable the measure (Stone, 1978). When creating a composite measure is not feasible, inter-rater reliability, if applicable, is another technique to assess reliability. Other techniques include test-retest reliability

Noise refers to extraneous variance related to the experimental setting – for example, smells, fatigue, noise and other distractions (Stone, 1978). When these extraneous variables affect participant performance in the experiment, measurement error increases. In order to reduce noise, the setting should be designed such that these distractions and factors affecting performance, other than the treatment condition, are minimized.

6.2.3 Research Design

Since one of the goals of this research was to test the propositions developed in Chapters 3 through 5, including those associated with the NKD design theory, the experiment was deemed an appropriate research design. The novelty of NKD tools implies that most organizations have not yet adopted this technology, further highlighting the appropriateness of this design (Benbasat, 1990). Due to the novelty of the technology and desire to maximize both internal and external validity, a laboratory and field experiment were conducted.

The NKD design theory was tested by evaluating the utility of a tool – the updated version of Athens – that adheres to the proposed NKD design principles. Utility is defined as the effectiveness and usefulness of a design in solving and addressing the stated problem – the meta-

requirements (Venable, 2006). The effectiveness of the NKD tool for discovering novel knowledge (level 3 knowledge-discovery task), the treatment condition, was tested against the effectiveness of tools at levels 1 and 2, the control and comparison conditions. Testing the effectiveness of a design solution in comparison to other solutions is an important aspect of evaluating a design theory (Venable, 2006). Furthermore, the effectiveness of different use processes was evaluated, focusing on measuring the radicalness and potential benefit of the chosen ideas to the organization.

The experimental design included two separate experiments: 1) lab experiment – to evaluate the effectiveness of the NKD tool, at the individual level of analysis, in comparison to the tools at the other levels of knowledge discovery, and 2) field experiment – to evaluate the processes of using an NKD tool in the multi-level organizational learning process. This is a valid approach for two reasons: 1) it allows us to maximize statistical power in the first experiment, and 2) it allows us to separate the effect of the tools on individual learning processes and discovery of novel knowledge, from the group interpretation and integration processes.

6.2.4 Lab Experiment

The lab experiment involved evaluating the effectiveness of the NKD tool, at the individual level of analysis, in comparison to the tools at the other levels of knowledge discovery, covering propositions 1 through 5. This experiment used a multiple equivalent control groups⁷, with pretest design. Student participants were used for this experiment. The experimental task involved asking each participant to search for the best ideas to solve a particular question or problem requiring novel knowledge – in this case, a knowledge strategy question.

Three treatment conditions, listed in Table 6.2, were used to compare the effect of knowledge-discovery tools on individual learning processes, learning modes and the discovery of novel knowledge. Student participants were assigned randomly to conditions, forming three

⁷ Participants were assigned randomly to groups, therefore the groups are considered “equivalent”.

treatment groups. Based on sample size calculations (see Appendix A), a minimum of 30 participants per treatment group were required.

Table 6.2 – Summary of Treatment Conditions

	Treatment Condition Description	Tool	Process	Tests
1	Level 1 knowledge-discovery tool used by individuals.	Google	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2a, P2d, P3a, P4a, P4b, P4c, P5
2	Level 2 knowledge-discovery tool used by individuals.	Vivisimo	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2b, P2e, P3b, P4a, P4b, P4c, P5
3	Level 3 knowledge-discovery tool used by individuals.	Athens 2.0	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2c, P2f, P3c, P4a, P4b, P4c, P5

The key dependent variables that were measured and compared in order to assess the effectiveness of the NKD design theory at the individual-level and address propositions 1 through 5 include: 1) type of knowledge discovered, 2) type of learning mode supported and 3) type of individual learning process supported. The perceived radicalness and potential benefit of ideas was also assessed and compared. During the pretest, individuals were asked to rate their current level of knowledge about the subject matter of the task. They were also asked to rate their level of knowledge of this subject matter during the posttest. The other dependent variables were only measured at posttest. A number of other potential confounding variables were measured during the pretest and posttest in order to control for potential threats to validity. A more detailed description of the lab experiment, including construct measurement and development, is provided in Chapter 7.

6.2.5 Field Experiment

The field experiment, a “quasi-experiment”, involved evaluating the processes of using an NKD tool in the multi-level organizational learning process, addressing propositions 6 and 7. In addition, this experiment also provided additional validation of proposition 2c, 2f, 3c and 5. This “quasi-experiment” used a within-subjects control, with pretest design. Organizational

participants were recruited from within a firm, selected according to predefined criteria such as industry, size and reputation for innovation. The organization site criteria are listed in Table 6.3. The experimental task involved asking each participant to search for the best ideas to solve a particular question or problem requiring novel knowledge – in this case, a knowledge strategy question.

Table 6.3 – Organization Selection Criteria

Criteria	Example
Organizations that value novel knowledge.	1) Organizations in a competitive environment where innovating is important. For example, pharmaceutical or high-tech industry. 2) Organizations that have a need to keep up with emerging consumer trends. For example, marketing organizations.
Feasibility of conducting the experiment.	Organizations with a strategic business unit that can make strategic decisions and that are within reasonable travelling distance from Kingston.
Organizational use of the Web.	Organizations that rely on the Web for business information.

Three treatment conditions, listed in Table 6.4, were used to test the processes of using an NKD tool and the associated outcomes. As described in Appendix A, one group of 6 to 12 participants from the local organization chosen to participate in this experiment was needed. This group participated in all three treatment conditions, creating within-subjects control groups. Although random assignment was not used, the groups being contrasted were equivalent since the same participants completed each treatment condition. Thus, selection bias was not an issue in either the lab or field experiment. The results from treatment conditions 3, 4, and 5 were compared to assess the effectiveness of each process, including individual, group and collaborative-filtering processes, thus addressing propositions 6 and 7.

Table 6.4 – Summary of Treatment Conditions

	Treatment Condition Description	Tool	Process	Tests
3	Level 3 knowledge-discovery tool used by individuals.	Athens 2.0	Each individual used the tool to find, evaluate and suggest the best idea to senior management.	P2c, P2f, P3c, P5, P6a, P6b, P6c, P7a, P7b
4	Level 3 knowledge-discovery tool used by individuals in conjunction with collaborative-filtering tool.	Athens 2.0	Results from each individual's search were collected, combined and rated by each individual. Ratings were fed into a collaborative-filtering tool that recommended the best idea to senior management.	P6a, P6b, P6c, P7a, P7b
5	Level 3 knowledge-discovery tool used by individuals in conjunction with GSS tool ⁸ .	Athens 2.0	Results from each individual's search were collected, combined and rated by each individual. Group collectively evaluated and suggested the best idea to senior management using GSS support structures to build consensus.	P6a, P6b, P6c, P7a, P7b

The key dependent variables that were measured and compared in order to assess the processes of using an NKD tool included: 1) perceived radicalness, 2) perceived potential benefit, 3) satisfaction with the process, and 4) confidence in the decision. The type of knowledge discovered, learning mode and individual learning processes supported were also assessed in order to further validate the findings from the lab experiment. During the pretest, individuals were asked to rate their current level of knowledge about the subject matter of the task. They were also asked to rate their level of knowledge of this subject matter during the posttest. The other dependent variables were only measured at posttest, after each process was completed. A number of other potential confounding variables were measured before and after the experiment in order to control for potential threats to validity.

The experimental tasks for the lab and field experiments were pretested repeatedly to adjust the technology, task, instructions, training, measurement and time limits. Table 6.5 provides a summary of the criteria used to evaluate the experimental design. Appendix B contains a summary of the treatment conditions and experimental task for the lab and field

⁸ GroupSystems was used as the GSS tool.

experiments. A detailed analysis of the threats to validity is provided in Appendix C. A more detailed description of the field experiment is provided in Chapter 9.

Table 6.5 – Criteria Used to Assess Experimental Design

Validity Type	Definition (Shadish et al., 2003)	Evaluation Criteria	This Study's Implementation
Internal	Validity of inferences about the causal relationship between the manipulated (independent) and outcome (dependent) variables.	Rule out major threats to internal validity. The major threats most applicable to MIS research area: temporal precedence, selection bias, lack of training, too much training.	Provided training to all participants. Three distinct treatment groups were used in lab experiment. Five independent individuals with industry experience evaluated the results to avoid bias. Inter-rater reliability was assessed.
Statistical Conclusion	Validity of inferences about the correlation between the manipulated variables and outcome variables.	Ensure that the sample size provides sufficient power and that the statistical assumptions, such as distributional assumptions and independence assumptions are not violated.	Based on sample size calculations (Friendly, 2006; Sturdivant, 2006; Tabachnick and Fidell, 2001), a power level of 80% and significance level of .05 (Cohen, 1988) can be achieved with a sample size of 30 individuals per treatment condition for the lab experiment. The field experiment used non-parametric statistical techniques. (see Appendix A)
External	Validity of inferences about whether the results of this study hold over variation in setting, persons, treatment variables and measurement variables. Generalizability of results.	Evaluate how representative your sample is of the target population. Evaluate how realistic the experimental setting and treatment was to determine generalizability to the real world.	Maximized realism and generalizability by using organizational participants in field experiment. Experimental task was pre-tested. Experimental task in the field experiment was chosen by the participating organization as an important problem requiring novel knowledge.
Construct	Validity of inferences from the operationalization of the construct to the intended definition of the construct.	Ensure that what the participants believed the manipulation (treatment) to be was indeed the intended manipulation. (Manipulation check)	After the lab and field experiments, participants were asked to describe what they believed they were asked to do and the characteristics of the tool (treatment) they used in their task.

6.3 Case Study

6.3.1 Overview

A case study is an idiographic and qualitative research method that develops a deep understanding of a phenomenon in its natural setting (Benbasat et al., 1987). Case research can be used for exploratory or confirmatory research. However, case research is particularly useful

for examining a phenomenon where little is known, understanding the complexities of processes in a natural setting, and documenting experiences in practice (Benbasat et al., 1987). Some of the key characteristics of case study research are: 1) conducted in a natural, field setting, 2) multiple methods of data collection are used, for example, interviews and documentation, 3) small sample size, 4) lack of intervention, controls or manipulation, and 5) results are analyzed using qualitative techniques.

Every research method has certain strengths and weaknesses. The weaknesses and flaws will ultimately affect and limit the conclusions that can be derived. In order to overcome these limitations, a variety of methods with counterbalancing strengths and weaknesses should be used to obtain corroborating evidence, an approach typically referred to as triangulation (Jick, 1979; Scandura and Williams, 2000). Experiments typically have high internal validity, but low realism and external validity. Case study research, on the other hand, has high realism, but weaker controls and generalizability (Scandura and Williams, 2000).

6.3.2 Research Design

The goal of the case study was to explore how an organization uses an NKD tool in practice. The newness of this technology implies that little is known about how organizations use tools to discover novel knowledge. Since only one organization has experimented with Athens 1.0, a tool that is more closely aligned to the NKD design principles than other existing tools, a single, revelatory case was used (Yin, 1994). In this case study, I addressed the following high-level questions:

1. What knowledge challenges are they faced with?
2. To what degree is the organization using the NKD tool?
3. How are they using the NKD tool within the organization?
4. What types of knowledge are they discovering with the NKD tool? And how important is this knowledge to the organization?

5. More specifically, what is the importance of discovering novel knowledge in their organization?
6. How do they discover/identify novel ideas (with or without the NKD tool)? To what degree do tools play a role in the discovery/learning process?
7. What other tools are currently being used to discover novel knowledge (for example, level 1 and level 2 tools)? How are they being used?
8. What challenges are they encountering in using the NKD tool? What NKD tool changes would they suggest? What NKD tool process/use changes would they suggest?

The interview questions were pre-tested prior to entering the organization and updated accordingly. Multiple sources of data were utilized: interviews with multiple organizational participants, observation of tool use, and review of search results from Athens, and other tools as applicable (see Table 6.6 for evaluation criteria). Interviews were recorded and transcribed. Each transcription was verified by the interviewee. All sources of data were entered into a qualitative analysis tool – NVivo – coded and analyzed. A detailed description of the case study procedure, data analysis and results is provided in Chapter 8. The design theory was updated based on the results of the case study. In addition, the results of the case study were compared and triangulated with the results of the experiments.

Table 6.6 – Criteria Used to Assess Case Study Design (Yin 1994)

Evaluation	Criteria	This Study's Implementation
Construct Validity	Multiple sources of evidence	Interviews – group and individual – and observation.
	Chain of evidence established	Coding and analysis of evidence in NVivo.
	Key informants review the results	Participants reviewed their own transcripts and the final results.
Internal Validity	Pattern matching used	Pattern matching within and between cases – observations of use and interviews.
External Validity	Use of analytic generalizability	Assessed degree to which the case analysis supported the theory developed in this dissertation. Triangulated with results from the experiments.
Reliability	Degree to which the procedures are documented and the artifacts used for analysis are organized in a way that is easily accessible to an external reviewer	The procedure was documented (see Chapter 8) and can be verified by referring to the documents stored within NVivo.

6.4 Summary

In order to address the research questions and test the research propositions described in Chapters 3 through 5, the research design included three studies, combining qualitative and quantitative methods. Two experiments – a lab and field experiment – were conducted to test the research propositions. In addition, a case study was conducted to explore novel-knowledge discovery in practice. Due to its exploratory nature, the case study was originally scheduled as the first study to be completed. However, due to timing constraints at the participating organization, the case study was conducted shortly after the lab experiment. The field study was completed after the case study.

In Chapters 7 through 9 additional methodological details, as well as the results of the three studies discussed above, are provided. Chapter 7 discusses the lab experiment (Study 1); Chapter 8 discusses the case study (Study 2); and Chapter 9 discusses the field experiment (Study 3).

CHAPTER 7

STUDY 1 – LAB EXPERIMENT

7.1 Purpose

The purpose of Study 1 – lab experiment – was to test the theoretical framework and competing theories developed in Chapter 3, as well as evaluate the design theory developed in Chapter 4. The NKD design theory instantiation – Athens – was updated to address the gaps identified in the evaluation of the initial prototype. Thus, another key objective of the lab experiment was to evaluate Athens 2.0 – the next step in the build-evaluate cycle of the NKD design theory.

The lab experiment involved evaluating the effectiveness of the NKD tool, at the individual level of analysis, in comparison to the tools at the other levels of knowledge discovery, covering propositions 1 through 5. These propositions differentiate knowledge-discovery levels and tools and how they support learning.

According to the NKD design theory and the theoretical framework (associated with propositions 1 through 3 and proposition 5) – referred to as Theory 1 – the research hypotheses are as follows:

Knowledge Discovery Type

H1a: The perceived level of **deep** knowledge discovered will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H2a: The perceived level of **broad** knowledge discovered will be higher for the Vivisimo tool group than the Google and Athens tool groups.

H3a: The perceived level of **novel** knowledge discovered will be higher for the Athens tool group than the Vivisimo and Google tool groups.

Learning Mode

H4a: The perceived level of mental model maintenance support (**MMM**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H5a: The perceived level of mental model tuning support (**TUN**) will be higher for the Vivisimo tool group than the Google and Athens tool groups.

H6a: The perceived level of mental model building support (**MMB**) will be higher for the Athens tool group than the Vivisimo and Google tool groups.

Learning Process

H7a: The perceived level of information foraging support (**FORAGE**) will be higher for the Athens tool group than the Vivisimo and Google tool groups.

H8a: The perceived level of intuiting support (**INTUIT**) will be higher for the Athens tool group than the Vivisimo and Google tool groups.

H9a: The perceived level of interpreting support (**INTERP**) will be higher for the Athens tool group than the Vivisimo and Google tool groups.

H10a: The perceived level of interpreting support (**INTERP**) will be higher for the Vivisimo tool group than the Google tool group.

Hypothesis 7a requires further discussion. Although all tools are proposed to support information foraging, I propose that Athens will provide greater support than the other tools for a novel-knowledge discovery task. In other words, Athens will perform better at supporting information foraging for novel-knowledge discovery.

The technological frames competing theory discussed in Chapter 3 – referred to as Theory 2 – suggests a different set of hypotheses:

Knowledge Discovery Type

H1b: The perceived level of **deep** knowledge discovered will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H2b: The perceived level of **broad** knowledge discovered will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H3b: The perceived level of **novel** knowledge discovered will be higher for the Google tool group than the Vivisimo and Athens tool groups.

Learning Mode

H4b: The perceived level of mental model maintenance support (**MMM**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H5b: The perceived level of mental model tuning support (**TUN**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H6b: The perceived level of mental model building support (**MMB**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

Learning Process

H7b: The perceived level of information foraging support (**FORAGE**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H8b: The perceived level of intuiting support (**INTUIT**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.

H9b: The perceived level of interpreting support (**INTERP**) will be higher for the Google tool group than the Vivisimo and Athens tool groups.¹

A third theory being compared in this study represents an integration or “synthesis” of Theories 1 and 2. These two theories assume a null hypothesis of: H_0 = no differences between tools. However, if technological frames influence the use of Web-search tools as proposed, the null hypothesis for an integrated theory is one in which there are group differences: H_0 = Google will outperform the other two tools. The hypotheses from Theory 1 are then tested against this null hypothesis.

In the following sections, I describe the research design employed for the lab experiment as well as an overview of the data analysis. The results of the experiment are summarized and the implications for Theory 1, Theory 2 and Theory 3 discussed.

7.2 Research Design

As described in Chapter 6, the lab experiment uses a multiple equivalent control groups with pretest design. In the following sections I describe the constructs measured in the lab experiment as well as measurement development. Pretesting activities and the components of the experiment itself are also discussed.

¹ H10b is omitted because it duplicates H9b.

7.2.1 Measures

Descriptions of the constructs measured in this experiment are provided below, including the independent variable, dependent variable, manipulation check, experimental controls and covariates.

7.2.1.1 Independent Variable

The independent variable – the treatment conditions – used in this experiment was the type of tool used to complete a novel-knowledge discovery task (level 3 task). The three types of tools used were Google (level 1 tool), Vivisimo (level 2 tool) and Athens (level 3 tool). Accordingly, three treatment conditions were used. Participants were assigned randomly to these conditions.

7.2.1.2 Dependent Variables

The primary dependent variables measured at posttest and compared between treatment groups include: 1) type of knowledge discovered, 2) type of learning mode supported and 3) type of individual learning process supported. As cognitive-based constructs, these variables were subjectively assessed and reported by the participant. Most of these constructs have not been measured in prior research. Thus, new measures were constructed based on a review of the relevant literatures and the theoretical framework developed in Chapter 3. The exception was learning mode. Both mental model maintenance and mental model building have been measured by Vandenbosch and Higgins (1996). Thus, I adopted the measures developed by Vandenbosch and Higgins (1996) for these two constructs, with some modifications to adjust the questions to the context of this study. All measures were perceptual measures.

In addition to measuring the types of knowledge discovery and learning support, the final product of the task was measured – attributes of the idea proposed to address the task. In the creativity literature, two dimensions have been used to measure creative products –novelty and utility or value (Amabile, 1983; Couger and Dengate, 1992; Newell et al., 1962). In addition to

measuring the creativity level of an idea or product, the “paradigm-relatedness” has also been measured (Hender et al., 2002; Nagasundaram and Bostrom, 1995). The creativity level measures the originality of the idea – whether the idea has been expressed before (Hender et al., 2002). Paradigm-relatedness measures the degree to which the idea is incremental and exists within the current paradigm, or whether it is radical and results in a paradigm shift (Nagasundaram and Bostrom, 1995).

In concordance with the literature and the outcomes measured in Study 3 (see Chapter 9 – Field Experiment), the two product-related constructs measured in this study were: 1) perceived radicalness, and 2) perceived benefit of the idea. For additional details regarding these constructs, please refer to Chapter 5, where both of these constructs are discussed in detail.

Most measures of product creativity are subjective in nature, relying on the judgment of a panel of experts (e.g. Amabile, 1983; Couger and Dengate, 1992; Ocker et al., 1996). Judgment is based upon assessing the product against an agreed upon definition of creativity. An objective analysis of product creativity is challenging since it requires a set of measurable and verifiable criteria (Amabile, 1983; Couger and Dengate, 1992). Both the identification of this criteria and the measurement of it are challenging. Accordingly, the perceived radicalness and potential benefit of ideas were assessed in this study by five independent evaluators after the completion of the experiment.

7.2.1.3 Manipulation Check

As discussed in Chapter 6, it is important to perform a manipulation check at the end of an experiment to ensure that what the participants believed the manipulation (treatment) to be was indeed the intended manipulation. A manipulation check helps strengthen construct validity. Two types of manipulation checks were included. The first manipulation check assessed whether individuals understood what they were being asked to do. In other words, I confirmed that participants understood they were performing a novel-knowledge discovery task (level 3 task). In

the second manipulation check, participants were asked about their perceptions of the tool's functionality and characteristics to assess the degree to which the participant perceived they were using the intended level 1, 2 or 3 tool (Google, Vivisimo or Athens).

7.2.1.4 Experimental Controls

In addition to the manipulation checks, three other types of experimental control variables were measured – satisfaction with the tool, effectiveness of training, and clarity of the instructions provided during the experiment. Measuring and comparing these experimental variables across the three treatment groups helped to rule out rival hypotheses that differences in these experimental variables caused, or at least partially contributed to, the effect.

7.2.1.5 Covariates

Similarly, a number of covariates were measured to control for their effect on the dependent variables. Table 7.1 provides a list of the covariates measured.

Table 7.1 – Summary of Covariate Variables

Covariate	Source
Initial MP3 knowledge	Developed
Tool familiarity	Developed
Personal innovativeness with IT (PIIT)	(Agarwal and Prasad, 1998; Lewis et al., 2003).
Task-specific computer self-efficacy (in particular, Web-specific self-efficacy)	Adapted (Hsu and Chiu, 2004)
General computer self-efficacy	Adapted (Compeau and Higgins, 1995)

During the pretest, individuals were asked to rate their current level of knowledge about the subject matter of the task, which in this case was MP3 technology. They were also asked to rate their level of knowledge of this subject matter during the posttest. This enabled an assessment of the degree to which the treatment increased their level of knowledge about the subject.

Three of the covariates noted in Table 7.1 have been measured before in the literature: Web-specific self-efficacy (task-specific computer self-efficacy), general computer self-efficacy, and personal innovativeness with IT (PIIT). General computer self-efficacy and task-specific computer self-efficacy have been shown to influence task performance (e.g. Agarwal et al., 2000;

Marakas et al., 1998). General computer self-efficacy represents “an individual’s judgment of efficacy across multiple computer application domains” (Marakas et al., 1998 p. 129), whereas task-specific computer self-efficacy refers to “an individual’s perception of efficacy in performing specific computer-related tasks within the domain of general computing” (Marakas et al., 1998 p. 128). The task in this case is related to searching on the Web. Thus, I adapted the Web-specific self-efficacy measurement instrument for use in the pretest (Hsu and Chiu, 2004). To measure general computer self-efficacy, I adapted the measurement instrument developed by Compeau and Higgins (1995).

Personal innovativeness with IT (PIIT) refers to “the willingness of an individual to try out any new information technology” (Agarwal and Prasad, 1998 p. 206). Given the novelty of some of the tools being used in this experiment (i.e., Athens) and the novelty of the task itself, which in this case also deals with technology, PIIT was deemed an important variable to measure. I adopted the measurement instrument developed by Agarwal and Prasad (1998).

In addition to the covariates mentioned above, several demographic variables were measured: gender, language, year in university program and age.

7.2.2 New Measurement Development

New scales were developed for the following constructs: knowledge-discovery types, learning processes, learning mode (mental model tuning), perceived radicalness of the idea and perceived benefit of the idea. The development of measurement scales involves multiple steps. The process used to develop and validate the new measurement scales followed the process described by Churchill (1979) and Moore and Benbasat (1991). Table 7.2 provides a summary of this process, described in detail below.

Table 7.2 – New Measurement Development Procedure

Procedure	Implementation
1. Specify domain and boundaries of the construct	Conducted a literature review in order to specify theoretical domain for each construct being developed.
2. Generate sample of items	Generated multiple items for each construct based on the literature review.

Procedure	Implementation
3. Scale development testing	Used card-sorting technique to assess construct validity (convergent and discriminant validities).
4. Purify measure	Updated scale items and repeated card sorting until convergent and discriminant validities satisfactory.
5. Collect data	Survey Instruments developed, pretested and administered during the lab experiment.
6. Assess reliability and validity	Construct validity assessed using factor analysis. Internal consistency reliability assessed using Cronbach's alpha.
7. Refine scale	Suggestions made for future instrument refinements.

The literatures reviewed and discussed in Chapters 2 through 5 led to the domain specification for each construct. For example, the radicalness construct (see Chapter 5) included new outputs, processes and technologies. Value was not included in the radicalness construct, but measured separately in the benefit construct.

Given the domain specification of each construct, the next step in the process involved generating multiple items for each construct. The objective of item creation was to ensure content validity (Moore and Benbasat, 1991), which reflects the degree to which the items in the scale reflect the content universe (Straub, 1989). Thus, items were created to reflect all aspects of the domain specification for each construct. For example, items reflecting new outputs, processes and technologies were created for the radicalness construct.

The objective of the scale development testing phase was to assess the construct validity of the measurement scales. After creating an initial set of items for each scale, construct validity was assessed using card sorting, a common technique used for the initial assessment of construct validity. Card sorting involves asking judges to sort items into construct categories. Construct validity is assessed by looking at the consistency with which an item is placed in a category, testing convergent validity (with the construct) and discriminant validity (from other constructs).

During the card-sorting exercise, convergent and discriminant validities were assessed by looking at the reliability of item classifications. Specifically, this was accomplished by measuring the consistency of item placement across the judges as well as number of items placed into the “target” construct (Moore and Benbasat, 1991). Items were discarded, added and reworded based

on feedback from judges, placement consistency and placement in the target construct. The card-sorting process, discussed in more detail below, also helped identify items that were confusing and ambiguous.

After multiple iterations of card sorting to refine and purify the measures, the entire measurement instrument was tested (Moore and Benbasat, 1991). After iterative pretesting, the measurement instruments were administered during the lab experiment. The data was collected and analyzed for construct validity – using factor analysis – and internal consistency reliability. Based on the results, suggestions were made for future refinements to the instruments.

Details regarding the card sorting procedure and results are provided below.

7.2.3 Card Sort

As noted above, items were generated for each of the following scales: knowledge-discovery types, learning processes, learning mode (mental model tuning)², perceived radicalness of the idea and perceived benefit of the idea. Based on the rules described by Jarvis et al (2003), the latter two constructs are formative and the others are reflective. For the formative constructs, reflective indicators were added to each scale to help overcome the issues with measuring formative constructs (Jarvis et al., 2003). Where practical, multiple items were created to measure each construct in order to strengthen both measurement reliability and construct validity (Churchill, 1979).

The card-sorting exercise was completed in two separate phases. In phase 1, judges were provided with the individual items to sort into construct categories, but were not provided with the names of the construct categories. Judges were asked to provide labels for the underlying constructs they identified. In the second phase of card sorting, judges were provided with the names for each of the constructs and asked to place items into these constructs. These phases are discussed in more detail below

² Learning modes mental model maintenance and mental model building were included in the card sorting exercise to assess discriminant validity.

7.2.3.1 Phase 1 Card Sort

Two rounds of card sorting were performed in the first phase, where no construct categories were provided. All items were printed on cue cards, which were randomly ordered. Each card was numbered so that a consistent random order was maintained across all card sorts. Initially, there were 68 cards and 21 target categories. Judges were provided a standard set of instructions. Judges were first asked to complete a trial sort on ten sample items that were unrelated to the constructs in the primary card sort. This procedure helped ensure the judges understood the process and helped sensitize and cue awareness to ambiguous or unclear items. Judges were provided with two additional categories to capture items that were ambiguous/confusing or that didn't appear to fit anywhere: "Does not make sense/ Are confusing" and "Does not fit any category".

7.2.3.1.1 Round 1

In the first round of card sorting, four judges participated. All judges in the card-sorting exercise were graduate students in management, ranging in areas of concentration – Marketing, MIS, Organizational Behaviour, Accounting, and Economics. The labels provided by the judges were mapped to the intended constructs based on the items included in the grouping and how the label related to the intended constructs' description. In many cases, several "intended categories" were grouped together in a higher-order construct. A summary of the matching category labels identified by the judges is provided in Appendix D.

Some participants viewed the constructs as higher-order constructs with dimensions. For example, one of the participants created the category "in what way did the tool impact how much I know about the subject", which included the construct items for deep, broad and novel knowledge.

Some constructs tended to be grouped with other related constructs. For example, one item from "ethical analysis" tended to load onto "compatibility", while another tended to load

onto “stakeholder analysis”. Similarly, “cost benefit analysis” tended to be grouped with “overall benefits”. “Protection of organization strengths” also tended to be grouped with compatibility. Compatibility seemed, therefore, to be a higher-order construct. In addition, participants seemed to identify compatibility as either compatibility with the firm’s intangible components or tangible components.

During this round, some items were placed into the “Does not make sense/ Are confusing” and “Does not fit any category” categories. These items were reworded accordingly. There was some cross loading of knowledge types, which resulted in some rewording to clarify the distinction. There was also some cross loading of learning modes, which resulted in some rewording. Intuit, one of the learning process constructs, had a number of problems. Intuit items rarely loaded together. Therefore, some rewording was done to alleviate this issue. There was also some cross loading between knowledge types, learning process and learning modes. This is to be expected since these constructs are related to one another. For example, novel knowledge, intuit and mental model building would be expected to be related; deep knowledge and mental model maintenance would tend to be related; broad knowledge and tuning would also tend to be related. In general, there appeared to be more discriminatory issues in the knowledge type, learning mode and learning process categories than in the categories related to radicalness and benefit of the ideas. Items were reworded (see Appendix D for a summary of the reworded items) and round 2 was initiated.

7.2.3.1.2 Round 2

In the second round of card sorting, two judges participated. Only one item was noted as having issues with loading onto a closely related construct. The item was reworded accordingly to help differentiate it from the related construct (see Appendix D for a summary of the reworded items).

Many of the groupings in this round were similar to those in round 1 (see Appendix D for a summary). Similar to round 1, higher-order constructs were developed by the judges, masking some of the dimensionality. For example, broader knowledge was grouped with novel knowledge, and deeper knowledge was grouped with broader knowledge. As noted in round 1, there were groupings of related constructs such as novel knowledge, intuit and mental model building. Thus, a decision was made to proceed to phase 2, where categories were provided for the participants.

Inter-rater reliabilities were calculated for phase 1 – rounds 1 and 2. The level of agreement between judges in categorizing items was calculated using Cohen's Kappa (Cohen, 1960). Scores greater than .65 are considered to be acceptable (Moore and Benbasat, 1991). In addition, the frequency with which items were placed in the target construct – item placement ratio – was calculated. There is no established level of acceptable frequency for item placement ratio. For phase 1, the Cohen's Kappa was 0.48 and the item placement ratio was 0.57. Although these scores were low, I felt that both Cohen's Kappa and item placement ratio would improve, which they did, once categories were provided to the judges in phase 2.

7.2.3.2 Phase 2 Card Sort

In phase 2, judges were provided with the target construct categories and asked to sort the cards into those categories. Similar to phase 1, judges were provided with “Does not make sense/ Are confusing” and “Does not fit any category” categories to help identify confusing and ambiguous items. There were 68 cards and 18 target categories. The reduction in categories resulted from combining categories into high-order constructs – for example, compatibility tangible and compatibility intangible – as suggested by the results in phase 1.

7.2.3.2.1 Round 3

In round 3, two judges participated. The overall item placement ratio (0.65) and Cohen's Kappa (0.56) were improved compared to phase 1, however a number of problems with item

loadings were detected. Similar to phase 1, there were cross loadings between the knowledge discovery, learning process and learning mode categories. The issues and associated resolutions are provided in Appendix D. After these changes were made, round 4 was initiated.

7.2.3.2.2 Round 4

In round 4, four judges participated. The overall item placement ratio (0.75) and Cohen's Kappa (0.65) were much improved. For the radicalness and benefit categories, the Cohen's Kappa was close to 1. There were still a number of cross loadings between the knowledge discovery, learning process and learning mode categories, resulting in a lower Cohen's Kappa for these categories. In order to determine which constructs were most problematic, I analyzed the knowledge discovery, learning process and learning mode construct loadings combined across rounds 3 and 4 to get a bigger sampling of respondents. If an item loaded into the appropriate category at least 4 times out of 6, it was deemed to be non-problematic. Otherwise, the item was examined to see if a rewording, reloading or removal was required (see Appendix D for a summary of the issues and their resolutions). After these changes were made, round 5 was initiated.

7.2.3.2.3 Round 5

In round 5, two judges participated. The focus in this round was on the knowledge discovery, learning process and learning mode constructs since this is where the reliability issues resided. The cross loadings across categories in previous rounds made it difficult to get an accurate assessment of item placement consistency and item wording clarity. To help alleviate this issue, I provided the judges with the cards and construct categories grouped together by high-level category (knowledge discovery, learning process and learning mode) and included a description of this high-level category. The overall item placement ratio (0.84) and Cohen's Kappa (0.64) were similar to or improvements over round 4. The problematic items identified in round 4 were reviewed. The final resolution for these items is summarized in Appendix D.

7.2.4 Survey Development

After five rounds of card sorting, the revised items and scales were used to develop the survey instruments. To shorten the survey and avoid participant fatigue, a few items validated in the card-sorting exercise were dropped. At least three items per construct were retained, with the only exception being the information foraging and search-term definition items (only two items per design). Appendix D provides a summary of the items dropped and retained in the survey.

7.2.5 Experiment Pretest

After completing the card-sorting exercise and refining the measurement scales, several pretest phases were conducted to review and refine all aspects of the experiment. See Table 7.3 for a summary of the pretest activities.

Table 7.3 – Summary of Pretest Activities

Activity	Resources
Review experimental design	Professor and graduate student with significant expertise in experimental design.
Review task	Professor with expertise in knowledge management and knowledge strategy tasks.
Review questionnaire (including task)	5 graduate students in management.
Review lab (QEDC ³) set up and experiment flow	2 QEDC facilitators and 2 thesis committee members.
Pretest experiment, including training	2 graduate students, of which 1 is an expert in experimental design.

7.2.6 Experiment Pilot Test

A small pilot test was conducted with students from the Queen's School of Business Research Subject Pool. The purpose of the pilot test was to further test all aspects of the experiment and ensure any problematic instructions or questions were identified and corrected. No major issues were detected, but small changes to the instructions were made. Nine students signed up to participate in the pilot test, but only six students actually participated. The number of participants in each treatment group was 1, 2, and 3 for Google, Vivisimo and Athens respectively.

³ Queen's Executive Decision Centre

7.2.7 Experiment Execution

7.2.7.1 Sample

Participants in the experiment were undergraduate Commerce students from the Research Subject Pool. Students who participated received 1.5 course marks, as stipulated by the Research Subject Pool guidelines. As an incentive, the student with the best idea was promised a \$100 gift certificate to a store of his or her choice.

Experimental sessions were run over the course of a month and a half in the winter term of 2007 in the Queen's Executive Decision Centre. Logistically, a maximum of eight students could participate in each session. Each participant completed the experiment individually. Treatment conditions were randomly assigned to each experimental session. The number of participants in each session ranged from one to eight. The sample included 99 students; sample size for each treatment condition was 31 for Google, 35 for Vivisimo, and 33 for Athens.

7.2.7.2 Pretest

The pretest was administered to participants via Web survey (see Appendix E for samples of all questionnaires used in the lab experiment). The pretest questionnaire contained a letter of information, informed consent and questions measuring all of the covariates noted above, with the exception of PIIT. PIIT was measured in the posttest to limit hypothesis guessing during the completion of the experimental task. In addition, a description of the experimental task was provided and participants were asked to provide two search terms, which could be composed of multiple words. Athens can take six to ten hours to complete processing, therefore the provision of search terms prior to the experiment allowed the researcher to run the queries for the Athens treatment group. To facilitate this and treat all participants equally, participants were asked to complete the survey at least 24 hours before their scheduled experiment session. Participants who did not complete the pretest before the experiment were not permitted to participate in the experiment. The pretest questionnaire was estimated to take 10 to 15 minutes to complete.

7.2.7.3 Training

Prior to the start of the experiment, the researcher placed the following materials at each computer terminal in the QEDC lab: a copy of the information letter, which was also provided electronically in the Web survey; a consent form, which included further details about video taping and the use of direct quotes; a sheet with the participant's name and the search terms provided. Thus, participants were assigned to specific computer terminals in the lab. On each computer terminal the researcher opened the search tool assigned to the session and the Web form used to capture ideas during the experimental task. For the Athens sessions, the participant's specific Athens search results were opened. Laptop lids were then closed so that participants were not distracted during the training phase of the experiment.

At the start of the session, participants were asked to review the information letter – if they had not done so in the Web survey – and review and complete the consent form. Prior to starting the experimental task, participants were provided with instructions on completing the task and training on the assigned tool via Powerpoint presentation. Supporting handouts for the task and tool⁴ were also provided. A detailed script was prepared and used by the researcher to ensure consistency of procedure, instructions and training across all sessions and treatments (see Appendix F for a copy of the script). The only difference in the script across treatment conditions was the tool training. Training took approximately 10 to 15 minutes to complete.

7.2.7.4 Task

Upon completing training, participants were provided with 30 to 40 minutes to complete the task – a novel-knowledge discovery task (level 3). The task involved a case scenario in which participants were told they work for a company that designs, develops and sells MP3 players – a subject that most undergraduate students are familiar with (see Appendix G for a full description of the task). Participants were asked to generate four to five ideas that they felt were novel ways

⁴ Tool instruction handouts were provided for Athens only.

of applying the company's MP3 expertise, for example, to develop new services, new products, or new markets. They were told to use their Web-search results provided by the assigned tool to help generate these ideas – either reuse, modify or create new ideas. Participants were asked to document their ideas, search terms used and the Web link that spurred this idea (i.e., the stimulus) in the Web form provided. For the Google and Vivisimo treatment groups, participants were asked to use the search terms they provided in the pretest Web survey for their first search, but told they could search as many times using as many different search terms as they like. Due to the length of processing time, Athens participants searched through their existing results.

7.2.7.5 Posttest

After completing the task, participants were asked to rate their ideas on the radicalness and potential benefit of the idea. In addition, they were asked to rank order their ideas by how well the idea addressed the task. Rating and ranking were completed using GroupSystems⁵, a group-support software tool that provides voting capabilities suitable for rating and ranking. This process helped individuals identify which idea they felt was the best idea to propose forward for evaluation by independent judges. Participants were then asked to complete a Web survey (see Appendix E), in which they would describe their idea – both the idea content and rationale – and answer a number of questions about their experience using the tool to complete the experimental task. The questionnaire included the knowledge-discovery type, learning mode, learning process and PIIT measurement-scale questions. In addition, questions about the manipulation (manipulation check), MP3 knowledge, and experimental control questions were asked. Lastly, participants were asked several qualitative questions including:

- 1) What features of this tool assisted you with today's task?
- 2) What challenges did you encounter using this tool?
- 3) If this tool could be improved, what are the most important features that should be added or fixed?

⁵ The researcher received training on how to use GroupSystems prior to the experiment.

The posttest lasted approximately 20 to 25 minutes. The entire experiment – from pretest to posttest – was approximately 90 minutes in duration.

7.3 Analysis

Responses from the pretest and posttest surveys were matched and analyzed to test and compare the hypotheses for Theories 1, 2 and 3. A preliminary analysis, described in more detail below, was conducted in which missing values, outliers and normality were assessed across all variables.

7.3.1 Preliminary Analysis

The steps undertaken as part of the preliminary data analysis are described below, beginning with a summary of the variables used in the data analysis (Table 7.4).

Table 7.4 – Summary of Variables

Construct	Sub-constructs	Description
Independent Variables		
Tools	Google	Level 1 tool
	Vivisimo	Level 2 tool
	Athens	Level 3 tool
Dependent Variables		
Knowledge Discovery Type	DEEP	Deep knowledge
	BROAD	Broad knowledge
	NOVEL	Novel knowledge
Learning Mode	MMM	Mental model maintenance
	TUN	Mental model tuning
	MMB	Mental model building
Learning Process	FORAGE	Information foraging
	STERM	Search-term definition
	INTERP	Interpreting processes
	INTUIT	Intuiting processes
MP32	MP32	Level of knowledge after treatment
Covariates		
MP31	MP31	Original level of knowledge
Tool Familiarity	TOOLFAM	Level of familiarity with tool
PIIT	TOTPIIT	Personal innovativeness with IT
WSE	TOTWSE	Web-specific self-efficacy
GCSE	TOTGCSE	General computer self-efficacy

Construct	Sub-constructs	Description
Gender	GENDER	Gender (M or F)
Language	LANGUAGE	First language (English or other)
Year in Pgm	YEAR	Year in university program (1, 2, 3, 4, other)
Age	AGE	Age in years
Experimental Control Variables		
Satisfaction	SAT	Satisfaction with the tool
Training	TRAIN	Effectiveness of training
Clear Instructions	CLEAR	Clarity of instructions
Manipulation Check		
Level 1	LEVEL1	Level 1 tool characteristics
Level 2	LEVEL2	Level 2 tool characteristics
Level 3	LEVEL3	Level 3 tool characteristics
Manipulation	MANIP	Manipulation check of the task

7.3.1.1 Missing Value Analysis

A missing value analysis was conducted in SPSS (see Appendix H for the detailed results). There were 10 instances of missing values – only one case had more than one missing value. In this particular case, all the missing values were related to the manipulation check for tool functionality. Means were substituted for missing values. During the assessment of normality, one case was detected where age was listed as zero. For this case, the mean age was substituted (19.9). The missing value analysis was rerun and no cases of missing values were found.

7.3.1.2 Outlier Analysis

In order to check for univariate outliers, variables were saved as Z scores (standardized values). The plots and extreme values analysis in SPSS were reviewed by tool group. At the tool group level, variables that had standardized values greater than 3.29 were noted as outliers (Tabachnick and Fidell, 2001) (see Appendix H for the detailed results). In addition, variables that had standardized values greater than 3.00 but less than 3.29 were noted as approaching outlier status. At this stage in the analysis, nothing was done to the variables noted as outliers. This analysis was repeated after the normality assessment and variable transformations.

7.3.1.3 Normality Assessment

The results were reviewed for normality by looking for significant skewness and kurtosis ($Z > 1.96$) (Tabachnick and Fidell, 2001) by tool group (see Appendix H for the detailed results). Variables that showed significant skewness or kurtosis in the normality analysis were transformed. Although the analysis for normality was performed at a tool, variable level, the transformation itself was performed on the entire variable (across all tools). The type of transformation performed on each variable was based on the guidelines described in Tabachnick and Fidell (2001): for negative skewness, the variable was reflected and the square root taken; for positive skewness, the square root was taken.

After this initial transformation, the skewness and kurtosis of the transformed variable was reassessed, by tool. If significant skewness or kurtosis was still present, the logarithm of the reflected variable (or variable in the case of original positive skewness) was taken (Tabachnick and Fidell, 2001). An assessment of skewness and kurtosis was always performed after each transformation. If the transformation made the skewness or kurtosis worse (for either the group with problems originally, or a group that previously had no problems), this particular transformation was not used. A summary of the final transformations can be found in Appendix H.

Once all of the transformations were completed, the outlier analysis was repeated to see if the transformations removed the significant outliers. The only significant outliers after the transformation were with Vivisimo tool familiarity. However, this was acceptable because there was little variation in these scores. Thus, any deviation from the standard response would show up as an outlier. For example, most people answered “1 – strongly disagree” regarding their familiarity with Vivisimo. Thus, individuals who answered “2-moderately disagree” or “3 – disagree” showed up as outliers. Thus, no action was required for these “outliers”.

In order to check for multivariate outliers, a regression was run using all dependent variables post transformation. The Mahalanobis distance was calculated as part of this process.

With 32 dependent variables (individual items), and a $p = .001$, the critical χ^2 is approximately 60. No multivariate outliers were detected (see Appendix H for the detailed results).

7.3.1.4 Exploratory Factor Analysis

Since multivariate analysis of covariance (MANCOVA), the chosen analysis method, is appropriate when dependent variables are moderately but not highly correlated (Tabachnick and Fidell, 2001), a correlation table of the dependent variables was created (see Appendix I for details). As expected, there were a number of significant correlations between dependent variables. Thus, the next step was to conduct an exploratory factor analysis (EFA) on the dependent variables to assess the degree to which they represented underlying latent constructs/factors. In addition, an EFA was conducted for the covariates PIIT, GCSE and WSE to evaluate convergent and discriminant validities. Following the creation of composite variables, scale reliabilities were assessed and deemed satisfactory. The results are presented in more detail below.

7.3.2 Data Analysis

Using the composite variables created during the preliminary data analysis, MANCOVAs were performed to assess the degree of support for Theories 1, 2 and 3. A secondary analysis was performed to assess differences in the amount of knowledge gained at the end of the experiment. After analyzing these results, several other quantitative tests were completed: assessing the manipulation check and other experimental controls. The qualitative data collected was also analyzed.

7.3.3 Idea Evaluation

After completing all experimental sessions, the best ideas identified by each participant (one per participant) were compiled and provided to five independent judges for evaluation. All five judges were either currently working in industry in a technology-related field, or had worked in industry. There were 99 ideas in total to be evaluated on radicalness and benefit criteria. Ideas

were reviewed and minor edits were made (see Appendix J for editing rules). In order to avoid fatigue, the ideas were divided randomly into five groups and five separate Web surveys were created. Each judge was asked to complete the surveys in a different order to avoid any order bias.

For each idea, judges were asked a total of six questions – three pertaining to radicalness of the idea and three pertaining to the potential benefit. The questions were based on those tested in the card-sorting exercise. The number of questions per idea was limited to six to balance the need for construct validity and measurement reliability with evaluator fatigue. In a study by Couger and Dengate (1992), the creativity of IS products was measured using a panel of expert judges who assessed each product as high, medium or low on novelty and utility dimensions. To assist the judges in evaluating the products, a set of guidelines for low, medium, high novelty and utility were developed and provided to the judges to anchor their evaluations. Similarly, for each of the radicalness and benefit questions, anchors were provided for the high, medium and low levels for each type of question (see Appendix D for the idea-evaluation survey template).

In case judges were not familiar with MP3 technology, a two-page MP3 summary sheet was developed and provided to the judges (see Appendix D). Reviewing this summary sheet was not required.

Prior to sending the idea evaluation questionnaire to the five judges, it was pretested by five individuals – four PhD students in management and one individual working in the high-tech industry.

The results of these analyses are presented below.

7.4 Results

7.4.1 Exploratory Factor Analysis

Three individual exploratory factor analyses (EFA) were performed for the dependent variables knowledge-discovery type, learning processes and learning mode. In addition, an integrated EFA, including all of these variables, was performed to assess convergent and discriminant validity. Lastly, an EFA was performed for the covariates PIIT, GCSE and WSE. Each of these analyses is described below.

7.4.1.1 Knowledge-Discovery Types EFA

Since some variables were transformed and some were not, variables were not in a common scale. Thus, Z scores were created post-transformation and used for the factor analysis. Both orthogonal (varimax in SPSS) and oblique (oblimin in SPSS) rotations were attempted. Using the orthogonal rotation, two factors were suggested based on eigenvalues above 1.00. However, there appeared to be a lot of cross loadings amongst items. Thus, oblique rotation was attempted.

Using oblique rotation, two factors were extracted. As shown in the factor correlation matrix, the two factors are highly correlated ($> .32$), warranting oblique rotation (Tabachnick and Fidell, 2001) (see Table 7.5). The pattern matrix depicts the contribution of the factor to the item. Factor 1 appears to contribute strongly to items Broad4, Broad5, Deep3, Novel1, Novel2, Novel3, and Novel4. Factor 2 contributes to items Broad3, Deep2 and Deep4. The loadings in the pattern matrix are between .563 and .798 and, thus, range from good (.55) to excellent ($>.71$) (Comrey and Lee, 1992) (see Table 7.6).

Table 7.5 – Factor Correlation Matrix: Knowledge-Discovery Types

Factor	1	2
1	1.000	-.494
2	-.494	1.000

Extraction Method: Principal Axis Factoring.
Rotation Method: Oblimin with Kaiser Normalization.

Table 7.6 – Pattern Matrix: Knowledge-Discovery Types

	Factor	
	1	2
Zscore(TBROAD3)		.615
Zscore(TBROAD4)	-.731	
Zscore(TBROAD5)	-.752	
Zscore(TDEEP2)		.798
Zscore(DEEP3)	.728	
Zscore(DEEP4)		-.706
Zscore(NOVEL1)	.623	
Zscore(NOVEL2)	.623	
Zscore(TNOVEL3)	-.634	
Zscore(NOVEL4)	.563	

Extraction Method: Principal Axis Factoring.
 Rotation Method: Oblimin with Kaiser Normalization.
 Loadings less than .30 were excluded.

Based on the pattern matrix (Table 7.6), the novel items appear to load together. However, they also load with some deep and broad items. The second factor appears to include deep and broad items. Interpretation is difficult because some variables were transformed. Thus, the analysis was repeated after all the variables were transformed. Since the variables now used the same scale, Z scores were not used in this second analysis. The results were very similar to the results from the initial oblique rotation. However, the negative signs were eliminated, making the analysis simpler (see Table 7.7).

Table 7.7 – Pattern Matrix: Knowledge-Discovery Types (All Transformed)

	Factor	
	1	2
TBROAD3		.619
TBROAD4	.707	
TBROAD5	.751	
TDEEP2		.807
TDEEP3	.739	
TDEEP4		.694
TNOVEL1	.645	
TNOVEL2	.625	
TNOVEL4	.576	
TNOVEL3	.640	

Extraction Method: Principal Axis Factoring.
Rotation Method: Oblimin with Kaiser Normalization.
Loadings less than .30 were excluded.

Ideally, three factors would allow us to adequately test the research hypotheses. The factor plot suggests some interesting structure within the first factor (see Appendix K). It appears there is some differentiation between the novel items and the Deep3, Broad5 and Broad4 items. It is plausible that Deep2, Deep4 and Broad3 represent deep knowledge and Deep3, Broad5 and Broad4 represent broad or more relational knowledge (see Table 7.8 below). Composite variables were created for the hypothesized constructs by calculating the mean across the items. In addition, composite variables were created for the constructs that emerged in this factor analysis (different broad and deep groupings). Both sets of composites were assessed and compared for scale reliability.

Table 7.8 – Alternative Factor Groupings for Knowledge-Discovery Type Composite Variables

Deep Knowledge (tied to what I already know)	
BROAD3	This tool helped me develop a broader understanding of a topic with which I am already familiar.
DEEP2	This tool helped me learn incremental details about what I already know.
DEEP4	This tool helped me learn more specific details about what I already know.
Broad Knowledge (expansion of knowledge on the subject)	
DEEP3	This tool helped me learn deeper knowledge about the subject.
BROAD4	This tool helped me learn relationships to other similar concepts.
BROAD5	The tool helped me expand my knowledge of this subject.
Novel Knowledge (new knowledge)	
NOVEL1	This tool helped me learn concepts that are new to me.
NOVEL2	The tool helped me learn new concepts that I had thought were unrelated to the subject.
NOVEL3	This tool helped me learn about new knowledge previously unknown to me.
NOVEL4	This tool helped me learn novel associations to the topic that surprised me.

7.4.1.2 Knowledge-Discovery Types: Scale Reliability

Scale reliability was assessed by calculating the internal consistency reliability using Cronbach’s alpha. Alpha was calculated for both factor structures of broad and deep constructs (see Table 7.9).

Table 7.9 – Summary Cronbach’s Alpha: Knowledge-Discovery Types

NOVEL	.747
BROAD	.692
DEEP	.729
DEEP2	.791
BROAD2	.839 (could rise to .894 if Broad4 was removed)

As shown, the Cronbach’s alphas were all above 0.7 (Nunnally, 1978), with the exception of broad, and are therefore satisfactory.

7.4.1.3 Learning Process EFA

As discovered during the factor analysis of knowledge-discovery types, interpretation is easier when all variables used in the analysis are in the same scale. Accordingly, all variables pertaining to learning processes were transformed. The transformed variables were tested for normality and all remained normal. Since one of the variables transformed earlier used a logarithmic transformation and the others used the square root transformation, the variables were converted to Z scores.

Two distinct factors appeared – one related to search terms, and the other one containing the intuiting, interpreting and foraging items (see Table 7.10). Thus, these three constructs are highly correlated. The distinction of these three constructs from choosing and refining search terms makes logical sense. These two factors do not appear highly correlated and have limited cross loading, thus oblique rotation was not done.

Table 7.10 – Rotated Factor Matrix: Learning Process

	Factor	
	1	2
Zscore(TFORAGE1a)	.690	
Zscore(TFORAGE2)	.655	
Zscore(TINTUIT3)	.547	
Zscore(TINTUIT4)	.720	
Zscore(TINTUIT1)	.811	
Zscore(TSTERM1)		.394
Zscore(TSTERM2)		.598
Zscore(TINTERP3)	.429	
Zscore(TINTERP1)	.465	.321
Zscore(TINTERP2)	.695	

Extraction Method: Principal Axis Factoring.
 Rotation Method: Varimax with Kaiser Normalization.
 Loadings less than .30 were excluded.

Ideally, four factors would allow us to adequately test the research hypotheses. Thus, a four factor solution was requested. Principal components analysis was used as the extraction method because principal axis factoring could not converge on a solution. Oblique rotation was

used because of the correlation between the four factors and cross loadings observed with orthogonal rotation.

The resulting four factor solution helps to bring out the four factors that were originally hypothesized (see Table 7.11). The loadings in the pattern matrix are between .464 and .936 and, thus, range from fair (.45) to excellent (>.71) (Comrey and Lee, 1992). Composite variables were created for STERM, FORAGE, INTUIT AND INTERP. Composites were created by taking the mean of the Z scores.

Table 7.11 – Pattern Matrix: Learning Process

	Component			
	1	2	3	4
Zscore(TFORAGE1a)	.817			
Zscore(TFORAGE2)	.801			
Zscore(TSTERM1)		.775		
Zscore(TSTERM2)		.751		
Zscore(TINTUIT3)			.936	
Zscore(TINTUIT4)			.832	
Zscore(TINTUIT1)	.379		.464	
Zscore(TINTERP3)				.925
Zscore(TINTERP1)				.621
Zscore(TINTERP2)	.783			

Extraction Method: Principal Component Analysis.
 Rotation Method: Oblimin with Kaiser Normalization.
 Loadings less than .30 were excluded.

7.4.1.4 Learning Process: Scale Reliability

Scale reliability was assessed by calculating the internal consistency reliability using Cronbach’s alpha (see Table 7.12).

Table 7.12 – Summary Cronbach’s Alpha: Learning Process

STERM	.361
FORAGE	.700
INTUIT	.795
INTERP	.632

Except for INTUIT and FORAGE, Cronbach’s alphas were lower than the 0.7 reliability threshold. Part of the reliability issue may be due to the differences in scores between tools. Thus the reliability analysis was performed at a tool level of analysis (see Table 7.13).

Table 7.13 – Summary Cronbach’s Alpha by Tool: Learning Process

	A	G	V
STERM	-.166	.482	.528
FORAGE	.645	.670	.491
INTUIT	.871	.776	.672
INTERP	.461	.710	.715

The Cronbach’s alphas at a tool level depict differences across tools, with both INTUIT and INTERP achieving alphas above .70 for at least two tool groups each. STERM and FORAGE alphas are likely low due to the low number of items included in these measures. FORAGE is close to the .70 threshold for acceptable reliability.

7.4.1.5 Learning Mode EFA

As discovered during the factor analysis of knowledge-discovery types, interpretation is easier when all variables used in the analysis are in the same scale. Accordingly, all variables pertaining to learning mode were transformed. Since one of the variables transformed earlier used a logarithmic transformation and the others used the square root transformation, the variables were converted to Z scores.

The transformed variables were tested for normality. There were two variables that had problems. MMB3 is now slightly negatively skewed and has positive kurtosis for Athens. After reviewing the histogram, it appears acceptable to keep the transformation despite the moderate skewness/kurtosis. Also, TUN5 is now very slightly negatively skewed for Google. Based on the improvements to interpretation in the factor analysis, these transformed variables were retained.

Two distinct factors appeared – one related to mental model building (MMB) and the other related to mental model maintenance and tuning (see Table 7.20). Thus, these two constructs are highly correlated.

Table 7.14 – Rotated Factor Matrix: Learning Mode

	Factor	
	1	2
Zscore(TMMM1)	.637	
Zscore(TMMM3)	.571	
Zscore(TMMM4a)	.669	
Zscore(TMMM5)	.739	
Zscore(TTUN2)	.624	
Zscore(TTUN3)	.636	
Zscore(TTUN4)	.733	
Zscore(TTUN5)	.777	.362
Zscore(TMMB3)		.578
Zscore(TMMB4)		.716
Zscore(TMMB6)		.665
Zscore(TMMB7)		.716

Extraction Method: Principal Axis Factoring.
 Rotation Method: Varimax with Kaiser Normalization.
 Loadings less than .30 were excluded.

Looking at the factor plot, it appears that tuning (TUN) and mental model maintenance (MMM) can be distinguished as separate factors (see Appendix K). As a final step, the factor analysis was rerun with a request for 3 factors (see Table 7.15).

Table 7.15 – Rotated Factor Matrix: Learning Mode (3 Factor)

	Factor		
	1	2	3
Zscore(TMMM1)	.450		.444
Zscore(TMMM3)			.731
Zscore(TMMM4a)	.418		.536
Zscore(TMMM5)	.495		.552
Zscore(TTUN2)	.513		.351
Zscore(TTUN3)	.691		
Zscore(TTUN4)	.871		
Zscore(TTUN5)	.628	.348	.451
Zscore(TMMB3)		.557	
Zscore(TMMB4)		.743	
Zscore(TMMB6)		.651	
Zscore(TMMB7)		.742	

Extraction Method: Principal Axis Factoring.
 Rotation Method: Varimax with Kaiser Normalization.

Reviewing the rotated factor matrix (Table 7.15), it appears that there is some evidence of separate tuning (TUN) and mental model maintenance (MMM) factors. However, the number of cross loadings between these two factors suggests they are highly related. The loadings in the factor matrix range from poor (.32) to excellent (>.71) (Comrey and Lee, 1992). Based on these results, and those of the card sort, composite factors were developed for the MMM, MMB and TUN constructs by calculating the mean of the Z scores

7.4.1.6 Learning Mode: Scale Reliability

Scale reliability was assessed by calculating the internal consistency reliability using Cronbach's alpha (see Table 7.16).

Table 7.16 – Summary Cronbach's Alpha: Learning Mode

MMM	.778
TUN	.839
MMB	.768

The Cronbach's alpha's are all well above the .70 threshold, indicating satisfactory reliability.

7.4.1.7 Full Factor Analysis

As a final step, an exploratory factor analysis containing all of the dependent variables was conducted. Using orthogonal rotation in the initial EFA, a number of cross loadings were found and the factors appeared to be highly correlated. Thus, oblique rotation was attempted. Seven factors emerged with eigenvalues over 1.0 (see Table 7.17)

Table 7.17 – Pattern Matrix: All Variables

	1	2	3	4	5	6	7
TBROAD3				0.564			
TBROAD4					0.702		
TBROAD5	0.686						
TDEEP2				0.800			
TDEEP3	0.606				0.367		
TDEEP4				0.632			
TNOVEL1	0.684						
TNOVEL2					0.632		
TNOVEL4					0.403		
TNOVEL3	0.512						
Zscore(TMMM1)				0.567			-0.406
Zscore(TMMM3)			0.327				-0.490
Zscore(TMMM4a)						0.309	-0.395
Zscore(TMMM5)				0.361			-0.392
Zscore(TTUN2)	0.416						
Zscore(TTUN3)	0.675						
Zscore(TTUN4)	0.711						
Zscore(TTUN5)	0.305						-0.347
Zscore(TMMB3)		0.661					
Zscore(TMMB4)		0.658					
Zscore(TMMB6)		0.672					
Zscore(TMMB7)		0.606					
Zscore(TFORAGE1a)				0.407		0.514	
Zscore(TFORAGE2)						0.641	
Zscore(TINTUIT3)							-0.347
Zscore(TINTUIT4)	0.473						
Zscore(TINTUIT1)						0.587	
Zscore(TSTERM1)			0.361				
Zscore(TSTERM2)			0.695				
Zscore(TINTERP3)						0.458	
Zscore(TINTERP1)						0.517	
Zscore(TINTERP2)						0.652	

Extraction Method: Principal Axis Factoring.
 Rotation Method: Oblimin with Kaiser Normalization.
 Loadings less than .30 were excluded.

Factor 1 appears to represent broad and novel knowledge as well as mental model tuning.

Two of the three items that emerged in the broad2 factor above (see alternative factor groupings

in Table 7.8) were represented in this factor. Interestingly, two of the novel items load onto factor 1 and the other two onto factor 5. The novel items that load onto factor 1 discuss new knowledge whereas the items that load onto factor 5 represent surprising new relationships. Thus, this factor analysis illuminates two different dimensions of novel knowledge.

Mental model building loads cleanly onto factor 2 and search-term definition loads onto factor 3.

Factor 4 appears to represent deep knowledge. The three items that loaded onto the emergent deep2 factor above (see alternative factor groupings in Table 7.8) loaded onto this factor.

In addition to the two “surprisingly” novel items, one of the broad items – Broad4 – loads onto factor 5. This particular broad item focuses on relationships. Thus, factor 5 appears to represent relationships – surprising and non-surprising relationships.

Factor 6 represents interpreting and information foraging. The combination of these items makes sense given the wording of the interpreting items – refined, summarized, meaningful results.

Lastly, factor 7 represents mental model maintenance.

Although there were some cross loadings, for the most part the loadings were clean and made sense. The only construct that did not load onto any of the factors in a meaningful way was intuiting.

Another EFA was performed, removing some of the constructs to simplify the analysis. Foraging and search-term definition items were removed since these constructs loaded cleanly and only have two items each and may be affecting the other factor loadings.

Eight factors were requested, since that is the number of constructs hypothesized to represent the items included in the analysis. The factors loadings are presented below (Table 7.18).

Table 7.18 – Pattern Matrix: All Variables (8 Factors)

	1	2	3	4	5	6	7	8
TBROAD3				0.5346				
TBROAD4					0.5398			
TBROAD5	0.7806							
TDEEP2				0.8549				
TDEEP3	0.5993							
TDEEP4				0.6171				
TNOVEL1	0.5976							
TNOVEL2					0.7594			
TNOVEL4								-0.4902
TNOVEL3	0.4532							
Zscore(TMMM1)				0.4732		-0.3933		
Zscore(TMMM3)						-0.6054		
Zscore(TMMM4a)						-0.4308		
Zscore(TMMM5)						-0.5437		
Zscore(TTUN2)	0.4066							
Zscore(TTUN3)	0.7475							
Zscore(TTUN4)	0.7887							
Zscore(TTUN5)	0.3566							
Zscore(TMMB3)		0.6406						
Zscore(TMMB4)		0.6774						
Zscore(TMMB6)		0.6791						
Zscore(TMMB7)		0.6424						
Zscore(TINTUIT3)						-0.3318		
Zscore(TINTUIT4)	0.5132							-0.3995
Zscore(TINTUIT1)								-0.5682
Zscore(TINTERP3)			0.8833					
Zscore(TINTERP1)							0.4638	
Zscore(TINTERP2)							0.7214	

Extraction Method: Principal Axis Factoring.
 Rotation Method: Oblimin with Kaiser Normalization.
 Loadings less than .30 were excluded.

Factors 1, 2, and 4 are the same as in the analysis above. Factor 5 still appears to represent relationships, however now only includes two items – Broad4 and Novel2. Factor 6 represents mental model building – similar to factor 7 in the previous analysis. Factor 7 represents interpreting, including items Interp1 and Interp2. Interp3 loaded onto factor 3, which

as a one item factor is less meaningful. Lastly, factor 8 represents surprising insights and includes one novel item (Novel4) and two intuiting items (Intuit1 and Intuit4).

There were fewer cross loadings in this final EFA and the intuit items loaded to a factor in a meaningful way. Although some hypothesized constructs loaded with other constructs to form meaningful higher-level constructs, overall both of these EFAs provide evidence of the convergent and discriminant validities of these constructs.

7.4.1.8 Summary: Scale Reliability

Scale reliability for the new factors identified above was assessed by calculating the internal consistency reliability using Cronbach’s alpha (see Table 7.19).

Table 7.19 – Summary Cronbach’s Alpha: New Factors

FACTOR 1 (BROAD, NOVEL, TUNING)	.871
BROAD3	.894
NOVEL3	.773
TUN	.839
RELN3	.722
INSIGHT4	.799
TINTERP3	.619

When compared to the scale reliabilities of the original factor compositions (see Table 7.20), it appears that there is minor improvement in some cases (for example BROAD3, NOVEL3, INSIGHT4), but a reduction in reliability in other cases (for example INTERP3, RELN3). The search-term definition construct was removed from subsequent analyses due to reliability issues and ambiguity in the item wording for the Athens treatment group. This issue is described in more detail below.

Table 7.20 – Summary Cronbach’s Alpha: Original Factors

NOVEL	.747
BROAD	.692
DEEP	.729
DEEP2	.791
BROAD2	.839 (could rise to .894 if Broad4 was removed)*
MMM	.778
TUN	.839
MMB	.768
STERM	.361
FORAGE	.700
INTUIT	.795
INTERP	.632

*see BROAD3 above

7.4.1.9 EFA Sample Size

There is a wide range of sample size recommendations for exploratory factor analysis (Costello and Osborne, 2005; Fabrigar et al., 1999; MacCallum et al., 1999). Some recommendations suggest using 5 to 10 participants per item measured, whereas others base their guidelines on the number of items measured per factor and the level of communality⁶. Across the three separate EFAs conducted in this study, there is a range of three (knowledge-discovery types, learning mode) to four (learning process) constructs hypothesized per area with three to four items being measured for each construct⁷. Given the maximum number of items measured in each area is 12, a sample size of 99 participants achieves a ratio of 8 participants per item. Further, given three to four items are measured per construct and wide communalities (ranging from .2 to .8), a sample size of 99 approaches the guideline of 100 participants (Fabrigar et al., 1999; MacCallum et al., 1999), but not more stringent guidelines – for example, minimum of 300 participants (Tabachnick and Fidell, 2001). Given the cross loadings and highly correlated

⁶ Communality refers to the amount of variance that is accounted for by the variable (Tabachnick and Fidell, 2001).

⁷ With the exception of FORAGE and STERM, as noted above.

factors identified in the analysis above, further work is recommended to further refine these measures for future research.

7.4.1.10 PIIT, GCSE and WSE EFA

An EFA including the constructs personal innovativeness with IT (PIIT), general computer self-efficacy (GCSE) and Web-specific self-efficacy (WSE) was performed to assess convergent and discriminant validities of these constructs. An EFA was performed for both the transformed and untransformed variables. There were negligible differences, therefore the untransformed results are shown below. There are some cross loadings between the GCSE and PIIT factors, but otherwise these items loaded cleanly onto the intended factors (see Table 7.21).

Table 7.21 – Rotated Factor Matrix: PIIT, GCSE, WSE

	Factor		
	1	2	3
PIIT1		.644	
PIIT2		.838	
PIIT3		-.636	
PIIT4		.828	
GCSE1	.641	.428	
GCSE2	.694	.337	
GCSE3	.884		
GCSE4	.540		
GCSE5	.591	.311	
GCSE6	.445		
WSE1			.729
WSE2			.695
WSE3			.696
WSE4			.756

Extraction Method: Principal Axis Factoring.
 Rotation Method: Varimax with Kaiser Normalization.
 Loadings less than .30 were excluded.

Note – item PIIT3 is negatively worded and was reverse coded for composite creation and the evaluation of scale reliability (discussed below). Composites for PIIT, GCSE and WSE were created by calculating the mean of the individual item scores. The transformed Z scores were used for the WSE composite. For the GCSE, only three of the six variables required

transformations. Thus, for measurement consistency the untransformed variables were used in the creation of the GCSE composite.

7.4.1.11 PIIT, GCSE and WSE: Scale Reliability

Scale reliabilities for PIIT, GCSE and WSE were assessed by calculating the internal consistency reliability using Cronbach's alpha (see Table 7.22).

Table 7.22 – Summary Cronbach's Alpha: PIIT, GCSE, WSE

PIIT	0.858
GCSE	0.844
WSE	0.833

The Cronbach's alpha's are all well above the .70 threshold, indicating satisfactory reliability.

7.4.1.12 Summary

As discussed above, the EFA for knowledge-discovery types suggested some minor underlying factor-composition differences for the deep and broad knowledge constructs from those hypothesized and validated during the card sort. The reliability analysis of the composites – for both the hypothesized and emergent factor compositions – showed higher reliability for the emergent factors. Thus, the emergent factors for the deep and broad knowledge constructs were adopted for the MANCOVA.

An integrated EFA including all dependent variables was conducted – combining knowledge-discovery types, learning modes, and learning process variables – in order to assess convergent and discriminant validities. Many of the factor loadings were similar to those found in the separate EFAs. However, there were some minor differences in the suggested factor composition. Some hypothesized constructs loaded with other constructs to form meaningful higher-level factors. The novel items were dispersed to form three factors: a novel-knowledge factor (items Novel1 and Novel3); a relationships factor (items Broad4 and Novel2); and a surprising-insights factor (items Novel4, Intuit1 and Intuit4). Further, the broad construct and

interpreting construct were both suggested as two-item factors as opposed to the hypothesized three-item factors. Despite these minor differences, the integrated EFA provided further support of the convergent and discriminant validities of the hypothesized constructs. Differences in scale reliabilities between the hypothesized constructs and the alternative factor structures were negligible.

Composites for the alternative factor structures noted above were created and analyzed in parallel with the original composites developed from the separate EFAs. There were some minor differences in the results, discussed below.

7.4.2 MANCOVA

In multivariate analysis of covariance (MANCOVA), the ideal covariates are those that are correlated with the dependent variables, not the independent variables. In addition, the covariates themselves should not be highly correlated since each covariate results in a loss of degrees of freedom (Tabachnick and Fidell, 2001). Thus, correlation matrices were created for the dependent variables and covariates (see Appendix L for the results).

There were several significant correlations between the dependent variables and the covariates. Of interest is the fact that TOTWSE, TOTGCSE, PIIT, MP31 and TOOLFAMG (tool familiarity with Google) are all significantly correlated. TOOLFAMG was not included in the final MANCOVA due to its correlation with the other covariates and because of its relationship to the treatment (independent variable) (Tabachnick and Fidell, 2001).

Multivariate tests of significance are reported below in Table 7.23. Tests of between-subjects effects are reported in Appendix L. The omnibus test shows there are significant differences between tool groups (α 0.001). Covariate gender (α 0.05) was significant. However, the only dependent variables with significant differences between tools were TOTDEEP2 (α 0.05), TOTBROAD2 (α 0.1), TOTFORAGE (α 0.001), and TOTINTERP (α 0.05).

Table 7.23 – Multivariate Tests of Significance

Effect	Wilks' Lambda	F	df	p
TMP31	0.897	0.915	10; 80	0.523
TOTPIIT2	0.861	1.289	10; 80	0.251
TOTGCSE3	0.879	1.104	10; 80	0.370
TOTWSE2	0.958	0.348	10; 80	0.964
AGE	0.947	0.447	10; 80	0.919
YEAR	0.937	0.541	10; 80	0.856
GENDER	0.797	2.034	10; 80	0.040
TOOL	0.562	2.670	20; 160	0.000

Given the significant differences between tool levels, pairwise comparisons were reviewed to assess support for Theories 1, 2 and 3 (see Appendix L for MANCOVA results). The null hypothesis used in the analysis of Theories 1 and 2 was: H_0 = no differences between tool groups. As discussed above, the null hypothesis used for Theory 3 was that Google will outperform the other two tools. Specifically, the following null hypothesis was used in the analysis: H_0 = Google will outperform the other two tools by 5 to 15%. This range was chosen based on its reasonableness. For Theory 3, the hypotheses from Theory 1 were tested against this null hypothesis. Support for each theory is summarized in Table 7.24 through 7.26 below.

Table 7.24 – Support for Theory 1 Research Hypotheses

H_0 – No differences between tool groups			
Knowledge Type			Significant findings?
H1a	Deep K	Google > Vivisimo, Athens	Yes at $\alpha = 0.1$ level
H2a	Broad K	Vivisimo > Google, Athens	No
H3a	Novel K	Athens > Vivisimo, Google	No
Learning Mode			
H4a	MMM	Google > Vivisimo, Athens	No
H5a	MTUN	Vivisimo > Google, Athens	No
H6a	MMB	Athens > Vivisimo, Google	No
Learning Process			
H7a	Forage	Athens > Vivisimo, Google	No
H8a	Intuit	Athens > Vivisimo, Google	No
H9a	Interp	Athens > Vivisimo, Google	No
H10a	Interp	Vivisimo > Google	No

Table 7.25 – Support for Theory 2 Research Hypotheses

H₀ – No differences between tool groups			
Knowledge Type			Significant findings?
H1b	Deep K	Google > Vivisimo, Athens	Yes at $\alpha = 0.1$ level
H2b	Broad K	Google > Vivisimo, Athens	No
H3b	Novel K	Google > Vivisimo, Athens	No
Learning Mode			
H4b	MMM	Google > Vivisimo, Athens	No
H5b	MTUN	Google > Vivisimo, Athens	No
H6b	MMB	Google > Vivisimo, Athens	No
Learning Process			
H7b	Forage	Google > Vivisimo, Athens	Yes at $\alpha = 0.01$ level
H8b	Intuit	Google > Vivisimo, Athens	No
H9b	Interp	Google > Vivisimo, Athens	Partially at $\alpha = 0.05$ (Google > Athens)

Table 7.26 – Support for Theory 3 Research Hypotheses

H₀ – Google will outperform other tools 5 – 15%				
	Construct	5%	10%	15%
Knowledge Type				
H1a	Deep K	accept	accept	accept
H2a	Broad K	accept	accept (A)/reject (V)	accept (A)/reject (V)
H3a	Novel K	accept	accept	reject
Learning Mode				
H4a	MMM	accept	accept	accept
H5a	MTUN	accept	accept	accept
H6a	MMB	accept	accept	accept
Learning Process				
H7a	Forage	accept	accept	accept
H8a	Intuit	accept	accept	accept
H9a	Interp	accept	accept	accept
H10a	Interp	accept	accept	accept

Only one hypothesis in Theory 1 is supported – H1a: The perceived level of deep knowledge discovered will be higher for the Google tool group than the Vivisimo and Athens tool groups. This hypothesis is also included in Theory 2 (H1b) and Theory 3.

The results for Theory 2 are stronger, with support for three hypotheses: H1b, H7b and H9b (partial). The adjusted group means for each dependent variable were reviewed. Google had

the highest mean score on all dependent variables. Thus, Theory 2 appears to have much stronger support than Theory 1.

The results for Theory 3 are also stronger than Theory 1. H1a is still supported. The results indicate that for the broad knowledge composite variable, the null hypothesis can be rejected for Vivisimo at the 10% threshold (Google will outperform by 10%). The fact that Vivisimo outperformed, or performed as well as Google, is significant since we expected Google to dominate across all outcomes. Thus, H2a is supported.

The null hypothesis can be rejected for the novel knowledge composite variable for Athens at the 15% threshold. Again, the fact that Athens outperformed Google is significant since we expected Google to dominate across all outcomes. Interestingly, the null hypothesis can also be rejected for Vivisimo at the 15% threshold. However, the Athens group mean score is higher than Vivisimo's for the novel knowledge composite variable. Thus, H3a is supported.

In parallel, a MANCOVA with the alternative composites suggested by the integrated EFA was performed. There were only a few minor differences from the results discussed above, summarized in the Table 7.27. Note that the null hypothesis can be rejected for the alternative composite variable "relationships" for Athens and Vivisimo at the 10% threshold. The relationship composite variable comprises one broad knowledge item related to relationships and one novel knowledge item related to surprising associations. Thus, the relationship composite variable covers both H2a and H3a. These results provide additional support for H3a and H2a.

Table 7.27 – Differences between MANCOVAs

Type of Analysis	Result with Original Composites	Result with Alternative Composites
Between-Subjects Effects	TOTINTERP significant at α 0.05 level.	TOTINTERP2 significant at α 0.01 level.
	TOTNOVEL not significant.	TOTNOVEL3 significant at α 0.1 level.
Adjusted Group Means	Google had highest mean on all dependent variables.	Google had highest mean on all dependent variables, with one exception: Vivisimo had highest mean for TOTRELN (relationships factor).
Pairwise Comparisons	Theory 2: H2b (Broad knowledge) not supported.	Theory 2: H2b (Broad knowledge) partially supported at α 0.1 level (Google > Athens).
	Theory 2: H9b (Interpreting) partially supported at α 0.05 level (Google > Athens).	Theory 2: H9b (Interpreting) partially supported at α 0.01 level (Google > Athens).
	Theory 3: H2a (Broad knowledge) null hypothesis rejected at 10% level.	Theory 3: H2a (Broad knowledge) null hypothesis rejected at 15% level.
	Theory 3: H3a (Novel knowledge) null hypothesis rejected at 15% level.	Theory 3: H3a (Novel knowledge) cannot reject null hypothesis.
	Not applicable to original composites.	Theory 3: H2a and H3a (Relationships) null hypothesis rejected at 10% level.

Qualitative results were reviewed to assess support for the competing theories. First, the responses to the task manipulation check were reviewed (see Table 7.28).

Table 7.28 – Qualitative Support: Task Manipulation Check

Group	Comment	Themes
Google	"I was trying to find ideas about novel applications to mp3 players. I was also trying to validate an idea that I already had , which was to create playlists on the go."	Searching for something specific. Validating an existing idea.
Vivisimo	"Any new idea that would be clear from what was seen on the results screen. This way I would not have to open each search result and actually read the page in depth. I wanted obvious answers in front of me , which I did not often get, hence why I only had 3 ideas. I am impatient and dislike reading webpages in more detail unless I know they will be 100% useful ."	Expectations of obvious "answers".
	"I was trying to find additional information and background facts about ideas that I already had a basic idea about . For 3 or 4 out of 5 of my ideas, I came up with a brief outline of what I was looking for in my head before I began searching. The search tool gave me the initial idea to search for hearing aid possibilities in MP3 players."	Searching for something specific. Validating an existing idea.
	"I was trying to generate new ideas for the company, but this was hard without specific knowledge of the direction I already wanted to go in , which I did not have."	Searching for something specific.
Athens	"I was attempting to find a novel idea for the application of MP3 technology with focus on wireless data streaming."	Searching for something specific.
	"Was trying to make links between an MP3 player and a potential service which could augment the MP3 experience like wireless headphones ."	Searching for something specific.

In the task manipulation check, participants were asked to describe what they were trying to accomplish in their search. While most participants adequately paraphrased the intent of the level 3 task, some participants' responses suggested behaviours consistent with Theory 2. These participants were searching for something specific, validating an idea they already had, and expecting obvious answers. These responses illuminate the expectations that participants had of their searching experience and the tool.

The qualitative results from the "tool challenges" question were also reviewed. Responses from the Athens treatment group were revealing (see Table 7.29)

Table 7.29 – Qualitative Support: Tool Challenges

Group	Comment	Themes
Athens	"Being new to Athens, I required a few minutes to play around on it to better familiarize myself with its setup. Also, I was a little overwhelmed with all of the information at the beginning, while I <u>simply found it different from typical search engines.</u> "	Comparisons to Google.
	"Did not give descriptions of website links <u>like traditional search websites (aka - Google).</u> "	Comparisons to Google
	"I found it almost <u>no better than a regular search tool, such as Google.</u> "	Comparisons to Google
	"It is difficult to say, because the tool uses a different approach. I would appreciate, if it <u>provided some of the Google-type search too</u> , because it would be more helpful with the search words I chose."	Comparisons to Google
	"My search terms were not relevant enough to <u>what I was looking for</u> so I came up with many clusters that were not relevant to <u>what I wanted to do.</u> "	Expectations of what they wanted to find
	"A lot of the results had little to do with mp3's or anything that I would have <u>wanted to find.</u> "	Expectations of what they wanted to find
	"The results were irrelevant to <u>what I was looking for.</u> "	Expectations of what they wanted to find
	"The searches were <u>not what I expected.</u> "	Expectations of what they wanted to find
	"This tool <u>does not clearly attack the queries asked of it.</u> Although it provides a broad spectrum of options, much of the time I found myself sifting through unnecessary websites that prolonged my search."	Expectations of what they wanted to find

When asked to describe the challenges they encountered while using the tool, several participants in the Athens treatment group compared the tool to Google. Changing the tool to make it more similar to Google was viewed as desirable. Further, some participants in the Athens treatment group expressed comments about their expectations of what they would find. The fact that Athens produced results that were different from what they expected was viewed as a

challenge. Similar to the qualitative results reviewed for the manipulation check, a subset of participants were looking for something specific.

Thus, the behaviours of some participants in each of the treatment groups resemble focused search behaviours (level 1 task), not novel-knowledge discovery behaviours where surprises are expected and desired. Participants found it challenging to break out of a level 1 task mindset to perform a level 3 task. In fact, it appears that some users were performing a level 1 task (focused search) rather than the intended level 3 task (novel-knowledge discovery).

The results from the pairwise comparisons and the qualitative analysis reviewed above provide evidence of Google’s dominance and the influence of technological frames on the task and use of the tool. This suggests that Theory 3 is a more appropriate way to analyze the results. Further quantitative analysis of Theory 3 was conducted in order to better understand these results. This analysis was performed initially at an item level to see if the composite variables were masking anything of interest. The results of these analyses at the individual item level are shown below in Tables 7.30 through 7.32.

Table 7.30 – Support for Theory 3 – Knowledge Discovery Types

H₀ – Google will outperform other tools 5 – 15%			
Item	5%	10%	15%
DEEP2	accept	accept	accept
DEEP3	accept	accept	accept (A)/ reject (V)
DEEP4	accept	accept	accept
BROAD3	accept	accept	accept
BROAD4	accept (A)/reject (V)	accept (A)/reject (V)	accept (A)/reject (V)
BROAD5	accept	accept	accept
NOVEL1	accept	accept	accept
NOVEL2	accept	accept	reject (A)/accept (V)
NOVEL3	accept	accept	accept
NOVEL4	accept	accept	reject

For knowledge-discovery types, the null hypothesis can be rejected for Athens for items Novel2 and Novel4 at the 15% threshold. Reviewing the wording of these two items in comparison to the other two novel items, it appears that Novel2 and Novel4 represent surprisingly

new knowledge; whereas Novel1 and Novel3 represent simply new knowledge. This provides partial support for H3a (2 out of 4 novel construct items supported).

The null hypothesis can be rejected for Vivisimo for items Broad4 (5% threshold) and Deep3 (15% threshold). Broad4 is focused on relationships and Deep3 is focused on deeper knowledge. Both Broad4 and Deep3 were part of the emergent factor representing broad knowledge. Thus, this analysis suggests that Vivisimo outperforms Google for these two aspects of broad knowledge – relationships and depth of knowledge – providing partial support for hypothesis H2a (2 out of 3 broad construct items supported).

Table 7.31 – Support for Theory 3 – Learning Modes

H₀ – Google will outperform other tools 5 – 15%			
Item	5%	10%	15%
MMM1	accept	accept	accept (A)/ reject (V)
MMM3	accept	accept	accept (A)/ reject (V)
MMM4	accept	accept (A)/reject (V)	reject
MMM5	accept	accept	accept (A)/ reject (V)
TUN2	accept	accept	reject (A) /accept (V)
TUN3	accept	accept	accept (A)/ reject (V)
TUN4	accept	accept	accept
TUN5	accept	accept	accept (A)/ reject (V)
MMB3	accept	accept	accept
MMB4	accept	accept	accept (A)/ reject (V)
MMB6	accept	accept	accept
MMB7	accept	accept	reject

For learning mode, the null hypotheses can be rejected for Vivisimo for all four mental model maintenance items at the 15% threshold and, in one case, the 10% threshold. This is somewhat surprising since Google was hypothesized to best support mental model maintenance. It is plausible that Vivisimo, although similar to Google, has some advanced features such as clustering that support mental model maintenance better than Google. Perhaps more surprising is the fact that the null hypotheses can be rejected for Athens for one mental model maintenance

item at the 15% threshold. Athens and Vivisimo both provide clustered results, which may provide higher levels of support for mental model maintenance. Thus, H4a is not supported for Theory 3.

The null hypothesis can be rejected for two of the mental model tuning items for Vivisimo at the 15% threshold – TUN3 and TUN5. These items refer to increasing and refining the individual's understanding of the topic. For TUN2, which refers to augmenting the individual's understanding of the topic, the null hypothesis can be rejected for Athens at the 15% threshold, but not for Vivisimo. While this is surprising, participants may have found that the broad scope and novelty of the Athens results helped augment their overall understanding. H5a is therefore partially supported (2 out of 4 mental model tuning construct items supported).

The null hypothesis for two of the mental model building items can be rejected – MMB7 and MMB4. In the case of MMB7, which refers to reframing the individual's thinking, the null hypothesis can be rejected for Athens at the 15% threshold. This is a very interesting finding given one of the novel-knowledge discovery design principles refers to reframing existing mental models. Thus, this result provides partial support for both the NKD design theory and H6a. The other item – MMB4 – refers to re-orienting the individual's thinking regarding the subject. The null hypothesis can be rejected for MMB4 for Vivisimo at the 15% threshold. This result is surprising, but provides evidence that level 2 tools may partially support mental model building. Thus, H6a is partially supported (1 out of 4 mental model building construct items).

Table 7.32 – Support for Theory 3 – Learning Processes

H₀ – Google will outperform other tools 5 – 15%			
Item	5%	10%	15%
FORAGE1	accept	accept	accept
FORAGE2	accept	accept	accept
INTUIT1	accept	accept	accept
INTUIT3	accept	accept	accept (A)/reject (V)
INTUIT4	accept	accept	accept
INTERP1	accept	accept	accept (A)/reject (V)
INTERP2	accept	accept	accept
INTERP3	accept	accept	accept

For learning processes, the null hypothesis can be rejected for INTUIT3 for Vivisimo at the 15% threshold. This finding, although counter to H8a, may suggest that Vivisimo’s clustering features and provision of metaknowledge help individuals develop insights from the results. In addition, the null hypothesis for INTERP1 can be rejected for Vivisimo at the 15% threshold. The wording of this item “the tool interpreted the raw results for me” suggests that clustering helps refine the results so that the user can more easily understand them. Thus, hypothesis H10a is partially supported (1 out of 3 interpreting construct items). However, hypotheses H7a, H8a, H9a are not supported.

7.4.2.1 Summary

A summary of the results for Theories 1, 2 and 3 is provided below in Table 7.33. This summary is based on the results from the pairwise comparisons and analysis of both the composite variables and individual variables. These results were also supported by a qualitative analysis of the manipulation check and tool challenges question. As shown in Table 7.33, Theory 3 received the strongest support.

Table 7.33 – Results Summary

Hypothesis	Theory 1	Theory 2	Theory 3
H1a/b (DEEP)	H1a supported	H1a supported	H1a supported
H2a/b (BROAD)	H2a not supported	H2b partially supported*	H2a supported
H3a/b (NOVEL)	H3a not supported	H3b not supported	H3a supported
H4a/b (MMM)	H4a not supported	H4b not supported	H4a not supported
H5a/b (TUN)	H5a not supported	H5b not supported	H5a partially supported
H6a/b (MMB)	H6a not supported	H6b not supported	H6a partially supported
H7a/b (FORAGE)	H7a not supported	H7b supported	H7a not supported
H8a/b (INTUIT)	H8a not supported	H8b not supported	H8a not supported
H9a/b (INTERP)	H9a not supported	H9b partially supported	H9a not supported
H10a (INTERP)	H10a not supported	NA	H10a partially supported

*For the alternative BROAD composite variable only

As a secondary analysis, an analysis of covariance (ANCOVA) was performed for the dependent variable representing the perceived level of MP3 knowledge after the treatment. The purpose of this analysis was to determine whether there were differences between tool treatment groups in the amount of knowledge gained about the subject – MP3 technology. There were significant differences between the Athens and Google tool groups (α 0.1) (see Appendix L for detailed results). Google participants reported learning higher levels of MP3 knowledge than Athens participants on average. This finding is congruent with the hypothesis that Google supports the discovery of specific, deep knowledge and mental model maintenance. The Athens tool should not have helped participants learn more about MP3 technology directly, but novel knowledge that could be associated with MP3 technology.

A review of the experimental controls analysis is provided in the following section.

7.4.3 Experimental Controls

Four different types of experimental controls were assessed based on the posttest data from the experiment: manipulation check, as well as differences in training, instruction clarity and tool satisfaction.

The manipulation check was comprised of a quantitative and qualitative component. For the quantitative component, the tool-functionality levels were analyzed using MANOVA to

determine whether each tool group perceived the appropriate tool characteristics for that level of tool. The omnibus test shows there are significant differences between tool groups (α 0.001).

The between-group differences, group means and pairwise comparisons were reviewed. Based on this assessment (see Appendix M for details) the manipulation was successful.

For the qualitative manipulation check, responses to the following question were reviewed for each tool group: “What were you trying to accomplish in your search?” If key words from each of the following groups of synonyms were present in the response, it was deemed acceptable: 1) new, novel, innovative, radical, unique or different; 2) ideas, products and services, ways, applications or uses. Otherwise, a careful analysis of the wording was done to assess whether the participant understood the purpose of the level 3 task. Most individuals seemed to understand the task. However, some individuals, as discussed above, had switched the task from a level 3 task to a level 1 task. A detailed analysis is provided in Appendix M.

For the final experimental control assessment, a MANOVA was conducted to assess any group differences in training, instruction clarity and tool satisfaction. The omnibus test indicated there were significant differences between tool groups (α 0.01). Reviewing the between-subjects effects, there were no significant differences between the tool groups for the training and clear instructions items. However, there were significant differences in the satisfaction scores. Specifically, there were significant differences between Athens and the other tool groups. Participants in the Athens tool group were clearly less satisfied with using Athens (see Appendix M for details). Given the novelty of the tool and the prototypical state of the tool, this is not surprising.

7.4.3.1 Summary

An analysis of the experimental controls provided evidence that the manipulation – tool level – worked successfully. Participants perceived the functionality of the assigned tool as intended. While most participants understood the task, some participants switched the task to a

level 1 task (focused search). There were no differences in the clarity of instructions and training across treatment groups. However, there were significant differences in the satisfaction scores. Participants in the Athens treatment group were less satisfied than the other treatment groups.

The qualitative results are discussed in detail below.

7.4.4 Qualitative Results

In the posttest questionnaire participants were asked three open-ended questions regarding their experience with using the assigned tool to complete the experimental task. These questions focused on the supporting features and challenges associated with using the tool and suggestions for improvement. The responses for each question were reviewed by tool group. Each response was assigned one or more theme(s)⁸, which described the response. The number of instances of each theme were tabulated for each tool group and then compared across tool groups. It is interesting to note that there were a number of themes that occurred across all three tool groups. These common themes, as well as any unique themes of note, are discussed below.

Analysis of the tool features question – “What features of this tool assisted you with today’s task?” – produced several themes that spanned either both the Athens and Vivisimo tool groups, or both the Vivisimo and Google tool groups. The most prevalent common theme between Athens and Vivisimo was the clustering feature, one of the tool characteristics proposed to be included in both level 2 and 3 tools. Themes common to Vivisimo and Google included advanced search features, basic search features and relevancy of results. The features identified by participants in each of these tool groups provided additional support for the tool-manipulation check – the features of each tool aligned with their hypothesized tool characteristics. Table 7.34 provides samples of the common themes for the features question. Further details can be found in Appendix N.

⁸ Themes were emergent, not predefined.

Table 7.34 – Features – Themes and Sample Comments

Tool	Sample Comments	Theme (# of instances)
Athens	"The clusters, which grouped relevant and similar information together and allowed me to drill down to find even more specific web links."	Clusters (18)
Vivisimo	"I liked the idea of clustering your search results into main headings. This allows me to locate my desired information faster and more effectively."	Clusters (18)
Vivisimo	"The advanced web search. I included a lot of the news sites because that's where new technologies would be featured."	Advanced search (4)
Google	"The advanced search form in Google helped me to get better results when I was looking for specific information."	Advanced search (5)
Vivisimo	"Search capability based on keywords I entered."	Keyword (basic) search (10)
Google	"I mainly used common keyword searches - nothing too complicated. I found that coming up with the right words and phrases to search was the key here."	Keyword (basic) search (6)
Vivisimo	"The way it presented results based on topics that most related to what I was searching for."	Relevant results (5)
Google	"The tool usually produces relevant results. It also usually tends to produce a wikipedia link, which was really useful for today's task."	Relevant results (5)

The comments about features within the Athens treatment group provided evidence that the tool adhered to the NKD design principles specified in Chapter 4. Table 7.35 provides a sample of the comments that indicate support for each design principle.

Table 7.35 – NKD Design Principles – Sample of Supporting Comments

Design Principle	Supporting comments
1. Resolve information overload problem by distilling to a small result set.	"The tool was extremely efficient at helping me find sites relevant to my searches, then further helping me drill down into clusters of them to find more specific and beneficial information." "The ability of the tool to group similar topics and keywords together was very helpful and reduced the time wasted by going through each link individually."
2. Resolve information overload problem by directing attention to most significant result set.	"The clustering and the key words at the top helped focus topics that are related to the search." "The clusters assisted me with today's task. The clusters made it easy to see where similar words were grouped and allowed me to quickly narrow down which cluster was more interesting."
3. Resolve information overload problem by reducing reliance on search term precision through clustering.	"I liked the clusters and how they grouped certain related keywords together. It allowed me to navigate into one, and narrow down my searches with further clusters." "The fact that the tool did not produce results that were specifically based on the words I entered. It sorted through the results to provide the distant connections rather than the websites that contained the words straight out."
4. Resolve uncertain relevance problem by providing metaknowledge for result set.	"Using the key words to assess the content of clusters and help me to decide whether to further investigate or not. The

Design Principle	Supporting comments
	summaries gave me a pretty good idea of what was being covered.”
5. Resolve uncertain relevance problem by providing measurement of “interestingness” to enable evaluation of novel knowledge.	“The tags feature seemed most useful. My particular search didn’t return very many results (maybe 20-30 pages) so I didn’t really need to use the sorting tools that much.” “The tool provided a cluster of websites that were potentially useful for me to complete the task. It also provided tags that sped up my searching process.”
6. Resolve difficulty learning novel concepts by reframing existing knowledge.	“Got me thinking outside of the box.” “The tool helped me brainstorm what types of services might be useful for my company, which do not come to mind right away.”
7. Resolve difficulty learning novel concepts by enabling the examination of existing mental models by juxtaposing existing knowledge with indirectly related stimuli or knowledge.	“It got information from web pages that connected MP3 technology to ideas I would never have thought of on my own.” “The ability to find concepts not directly connected to the search terms exposed me to potential applications for mp3 players I wouldn’t have otherwise thought of.” “Suggested ideas/terms that I didn’t really associate to MP3 initially.”

In addition, the following comments (see Table 7.36) exemplify how Athens supports the discovery of novel knowledge.

Table 7.36 – Novel Knowledge – Supporting Comments

Athens Tool Group – Sample of Responses to the Features Question
“Ability to find information that would not necessarily pop up using a normal search engine. E.g. user blogs with intense interest in mp3 tech.”
“It was helpful that the tool didn’t just provide any related information to my search words, but actually took those words and searched for novel ideas related to the subject matter. For example, if I were to search “digital music” on another search engine, it probably would have provided me with the actual music.”
“The fact that the tool did not produce results that were specifically based on the words I entered. It sorted through the results to provide the distant connections rather than the websites that contained the words straight out.”
“Some of the findings the tool provided were the least I would expect to draw my attention to otherwise.”
“Suggested ideas/terms that I didn’t really associate to MP3 initially.”
“It got information from web pages that connected MP3 technology to ideas I would never have thought of on my own.”
“The ability to find concepts not directly connected to the search terms exposed me to potential applications for mp3 players I wouldn’t have otherwise thought of.”
“The tool helped me brainstorm what types of services might be useful for my company, which do not come to mind right away.”

Comments regarding support for novel-knowledge discovery were unique to the Athens tool group. Thus, the qualitative analysis of the features responses provides evidence that the Athens tool provided the best support for the NKD design theory.

The analysis of the challenges question revealed common themes across two and sometimes all three tool groups. Both the Athens and the Google tool groups voiced their concern over irrelevant results. This was the most prevalent comment for both tool groups. One of the comments within the Athens tool group (see Table 7.37 below) provided further evidence that participants expected the results from Athens to be similar to Google – providing results that are directly connected to and related to the search terms. The comment – “This tool does not clearly attack the queries asked of it” – implies that, similar to Google, the Athens tool will do what you ask of it and assumes you are asking something specific in the first place. This provides further evidence regarding the existence and influence of the Google technological frame.

Further, a common theme found between both the Athens and Vivisimo tool groups was the comparison to Google and other search tools. The comparison to Google, in particular, provided support that Google is the technological frame for search tools.

Both the Athens and Google tool groups had comments related to the lack of metadata. However, these comments were not particularly prevalent.

There were two common themes that were both prevalent and spanned all three tool groups: challenges with search-term definition and refinement, and tool didn't give the individual what he or she was looking for. The former challenge supports the importance and difficulties associated with defining appropriate search terms for level 1 and 2 tools. Athens tool group participants were frustrated they couldn't change their search terms – something they are used to doing with tools like Google. However, as the comment for the Athens tool group below suggests, users had expectations based on their technological frames that the results would be related to their search terms. It was made very clear during the Athens training session that the results would not be related to the search terms.

The other common theme – didn’t give me what I wanted/ was looking for – mirrored the findings from the review of the manipulation-check question. Across all three tool groups, some participants were looking for something specific rather than scanning the results in an attempt to generate or reuse novel ideas. Again, individuals seemed to be approaching the task using a focused-search approach rather than a true level 3 task. Comments from the Google tool group were quite revealing:

“I found it harder to find interesting ideas by searching general keywords or phrases.”

“Sometimes it gave only the obvious results rather than the different ones.”

These comments imply that it was difficult to complete a level 3 task with a level 1 tool.

Table 7.37 – Challenges – Themes and Sample Comments

Tool	Sample Comments	Theme (# of instances)
Athens	“This tool does not clearly attack the queries asked of it. Although it provides a broad spectrum of options, much of the time I found myself sifting through unnecessary websites that prolonged my search.”	Irrelevant results (13)
Google	“It would often retrieve information/websites that were not relevant to my task. Sometimes it would display results that contained only one of the search terms instead of all of the ones I had used.”	Irrelevant results (8)
Athens	“I found it almost no better than a regular search tool, such as Google.”	Comparison to Google and other tools (3)
Vivisimo	“I didn’t like the fact that they didn’t provide me with spelling corrections, unlike Google if I mistyped a word they just searched for it as was and gave me useless results. I also didn’t like that you had to already know in what direction you wanted to search before starting.”	Comparison to Google and other tools (4)
Athens	“Did not give descriptions of website links like traditional search websites (aka - Google) -- resulted in clicking on links that were relevant to the topic.”	Lack of metadata (3)
Google	“The summaries were sometimes misleading. Once, the summary gave a relevant idea, yet when I clicked the link, I was directed to a list of articles. I did not know which article related to the summary I was interested in.”	Lack of metadata (1)
Athens	“The searches were not what I expected. Because my search words were less-focused, much of results were useless.”	Search-term definition and refinement (4)
Vivisimo	“I also didn’t like that you had to already know in what direction you wanted to search before starting.”	Search-term definition and refinement (3)
Google	“I had to continue to refine my search terms, as it is not always easy to figure out which words will produce the best results.”	Search-term definition and refinement (6)
Athens	“The results were irrelevant to what I was looking for.”	Didn’t give me what I wanted/ was looking

Tool	Sample Comments	Theme (# of instances)
		for (5)
Vivisimo	“The tool took the keyword but not what i was looking for always. It usually gave me sites that offered Mp3's for sale. I never indicated purchase at all in the search. I wanted information.”	Didn't give me what I wanted/ was looking for (13)
Google	“I found it harder to find interesting ideas by searching general keywords or phrases.”	Didn't give me what I wanted/ was looking for (5)

In addition to the common themes noted above, a number of Athens issues were highlighted in the “challenges” comments. Specifically, some of the comments made reference to the fact that the tool provided empty clusters. This was due to an error in Athens processing. Users could easily work around this issue, but the error was confusing and did not instill confidence in the tool. Further, a couple of comments provided evidence that Athens was overly reliant on proper names such as names of people and geographical places. In these cases, a portion of the tool results were skewed in a direction related to those names, providing irrelevant results. Lastly, participants found navigating through the tool and the organization of results confusing. Part of this was due to the novelty of the tool. However, the comments provided helpful suggestions to improve the usability of the tool.

Additional suggestions for improvements were provided by the responses to the suggested improvements question. Similar to the other two questions there was a common theme across Athens and Vivisimo tool groups in which comparisons to Google were made (see Table 7.38 below). There were several common themes detected across all three tool groups. These themes include improving the ability to filter and sort the results, providing additional metadata and quality indicators, improving navigation and providing support for search-term definition and refinement. It is interesting that in the filtering and sorting theme, participants from both the Google and Athens groups made the suggestion to enable the user to choose the degree to which results are related to the search terms. It appears that both participants made this suggestion

because they perceived the results from their respective tool to be indirectly related to their search terms, albeit by different degrees.

Another interesting finding was that some of the comments in the Athens tool group specifically suggested that the results be less novel and more related to the search terms entered.

For example:

“The common words found should be in close proximity to the words that were originally searched for.”

Clearly, Athens participants were not comfortable with the novelty of the results. However, some participants had suggestions for how to resolve this difficulty:

“I think some sort of function that would show you the connection to your original search term might help users make the connection”

“The tool should result in both direct associations as well as indirect associations.”

While the results of the qualitative analysis provided evidence that Athens has improved in its adherence to the NKD design principles, there is still a lot of room for improvement.

Table 7.38 – Suggested Improvements – Themes and Sample Comments

Tool	Sample Comments	Theme (# of instances)
Athens	“It is difficult to say, because the tool uses a different approach. I would appreciate, if it provided some of the Google-type search too, because it would be more helpful with the search words I chose.”	Comparison to Google (2)
Vivisimo	“Sponsored links provided at the side (similar to Google).”	Comparison to Google (1)
Athens	“Option to filter sources. Be able to choose the level of how abstract / indirect the results are.”	Filtering and sorting (4)
Vivisimo	“It would be good if there was a way to specify which kind of sites you were looking for - i.e. academic, or sales, or organization.”	Filtering and sorting (3)
Google	“In order to improve this tool, Google should allow people to decide whether they want results that either explicitly apply or loosely apply to the search words entered in order to decrease the number of unrelated pages that appear.”	Filtering and sorting (3)
Athens	“Including descriptions of website links so that you don't have to click on them to understand their content.”	Metadata (3)
Vivisimo	“Summary of what is in each site (which comes below the link).”	Metadata (1)
Google	“Potentially Google could provide more information on the hit (i.e. it comes from a news cite, Magazine article, etc.) along with its 2 line summary of the site that it shows under the web address link.”	Metadata (3)

Tool	Sample Comments	Theme (# of instances)
Athens	"I think it would be more useful to have the information better organized and categorized in descriptive English (as opposed to terms that one who hasn't been trained might not understand). The word 'cluster' was used for more than one level, and was already somewhat confusing."	Navigation (16)
Vivisimo	"I didn't really like the layout of the tool. I would also allow users to do an advanced search without having to first type in a search. For example, I don't think I could do an advanced search for the first search terms I tried."	Navigation (3)
Google	"More intriguing interface."	Navigation (1)
Athens	"More statistics would be better (e.g. relevancy, number of results / uniqueness)."	Quality indicators (2)
Vivisimo	"Ranking of page would be helpful."	Quality indicators (1)
Google	"Maybe it could show the percentage correlation per link? But I don't really have any issues with the tool."	Quality indicators (2)
Athens	"It should be easier to refine your search. Most people don't have a few hours to wait to search for ideas."	Search term refinement (6)
Vivisimo	"Possibly a thesaurus type program where if you type in a word, it gives you other similar words to try. For example, when I type in "portable music + uses" if a list came up somewhere that showed other possible synonyms like "mp3 player, walkman" for the first word, and "application" for the second word."	Search term refinement (3)
Google	"The most beneficial would be search within a search option. So let's say I got a result I liked, but the other results were poor. You should be able to click something like "similar to this result" which would help you target your search more fluently and efficiently."	Search term refinement (3)

7.4.4.1 Summary

An analysis of the qualitative data collected during the posttest revealed several common themes across tool groups, as well as interesting tool-specific themes. The qualitative results provided additional support for many of the quantitative findings, for example, the manipulation check and Google technological frame. The tendency of some participants to change the task to a level 1 task was evident. However, it was also clear that it was difficult for users to perform a level 3 task using a level 1 (Google) tool. Support was also provided for the adherence of Athens to the NKD design principles. Comments from the Athens group suggest that improvements to the tool are still needed.

7.4.5 Ideas

Prior to analyzing the idea-evaluation data, a preliminary analysis was conducted, including an assessment of inter-rater reliability (see Appendix O for the details). Overall, there was very low consistency in scores across the five judges. Consistency was lower for the benefit-related items than the radical-related items. However, this was also a problem in another study assessing creative products (Ocker et al., 1996).

For the main analysis, the mean score for each idea by question type (item) was calculated across the five judges. For each question type, analysis of variance (ANOVA) was performed to assess whether there were differences in the mean idea scores between the three tool groups. There were no significant differences. However, it is interesting to note that for the radical-related questions, Athens had consistently higher mean scores.

Composite scores were developed to create a total radical, benefit and overall idea score (see Appendix O for the details). There were no significant differences, but again Athens had a higher mean score for the total radical and overall scores.

7.5 Discussion

A series of competing hypotheses associated with the knowledge-discovery theoretical framework and technological frames were evaluated in this study. In addition, the NKD design theory and updated version of Athens were assessed. In a lab experiment three tools, each representing a different level in the theoretical framework, were compared in their performance of a level 3 task – novel-knowledge discovery. Multiple methods were used to test the research hypotheses. For example, self-report data was collected during the experiment using Web surveys, including closed-ended and open-ended questions. Both statistical analysis and qualitative analysis methods were employed. In addition to the self-report data, the ideas produced by participants during the experiment were subjectively evaluated by five independent

judges. The results of this study and the implications for novel-knowledge discovery are discussed below.

Comparing multiple competing theories provided rich insights into novel-knowledge discovery and the support that different types of tools provide. In this study, the hypotheses associated with three different theories were compared. Theory 1 represents the theoretical framework; Theory 2 represents the technological frames competing theory; the third theory represents the integration of Theories 1 and 2. The integration of these two theories involved evaluating the research hypotheses of Theory 1 using a null hypothesis that Google will outperform the other tools. When Theory 1, 2 and 3 were compared directly, Theory 3 received more support than either Theory 1 or Theory 2 alone.

Theory 3 provided support for all three of the knowledge-discovery type hypotheses. The hypothesis that Athens is more effective at supporting the discovery of novel knowledge than other tools was supported. Further, the analysis revealed that the novel construct may consist of multiple dimensions, one of these being surprising novel knowledge. Athens appears to be particularly effective at discovering surprising novel knowledge. Vivisimo was most effective at supporting the discovery of broad knowledge and Google the discovery of deep knowledge.

For the learning-mode hypotheses, Theory 3 provided partial support for two of the learning-mode constructs. The hypothesis that Google is most effective at supporting mental model maintenance was not supported. In contrast, both Athens and Vivisimo provided more support for mental model maintenance. The conjecture is that the clustering features within each of these tools provided better support for mental model maintenance than Google. The hypothesis that Vivisimo is most effective at supporting mental model tuning was partially supported. Interestingly, Athens also provided some level of support for mental model tuning. Again, the clustering functionality of both Athens and Vivisimo could be contributing to this finding. Lastly, the hypothesis that Athens provides the most effective support for mental model building was partially supported. The results suggest that Athens was most effective at reframing mental

models – a dimension of the mental model building construct. Vivisimo also provided some level of support for mental model building in terms of reorienting thinking.

While the support for the hypotheses associated with learning processes was weak, some interesting results emerged. Of the four hypotheses, only one received support (partial). The hypothesis that Vivisimo would be more effective than Google in supporting interpreting processes was partially supported in Theory 3. However, more interesting is the fact that Vivisimo, not Athens, outperformed Google on one of the intuiting items regarding the generation of insights. It is possible that the clustering feature in Vivisimo helped individuals see patterns and generate initial insights into the search results. Since Athens also has clustering functionality, this result may suggest that through additional refinement of the tool and training, Athens can also support intuiting effectively. The novelty of the Athens results, as well as the novelty of and usability issues within the tool itself may have limited its ability to support intuiting and interpreting processes.

The hypothesis that Athens would be more effective in supporting information foraging for a level 3 task was not supported in Theory 3. Instead, Google supported information foraging significantly better than both Athens and Vivisimo, thus, supporting the hypothesis associated with Theory 2. This finding can be explained by the qualitative findings, discussed below.

The qualitative analysis provided support for the NKD design theory and helped provide explanations for many of the quantitative findings. Specifically, the analysis provided support for the proposition that Google is the dominant design and has a strong influence on individuals' technological frames regarding Web-searching tools. Some participants in the Vivisimo and Athens tool groups compared these tools to Google in their written comments. Some participants felt that if each of these tools was more like Google, they would be “better” tools. In addition, some participants had difficulty changing their mindset to a level 3 task. Instead, some participants converted the task to a level 1 task where they looked for and expected to find specific results. Since Google, the dominant technological frame, is typically used for this type

of task, it makes sense that individuals would be more familiar with a level 1 task and that Google would perform better – for example, in information foraging. Thus, the Google technological frame has a strong influence on the individuals’ perceptions and expectations associated with Web-searching tools, which affected both cognition and behaviours.

Insight into the behaviours associated with search-term definition and refinement were found in the qualitative results. Participants found this process challenging, especially when searching for novel ideas. Athens participants were frustrated that they couldn’t change their search terms during the experiment. In fact, the quantitative results for the search-term definition construct (STERM construct) were excluded from the analysis above because of ambiguity for Athens participants⁹. Athens participants could answer the questions positively – “Yes I needed to refine my search terms” – or negatively – “I wasn’t permitted to refine my search terms”. However, the qualitative results helped provide insights into participants’ perceptions and behaviours regarding definition and refinement of search terms.

The qualitative results provided evidence that the updated Athens tool adheres more closely to the proposed NKD design principles than the initial prototype assessed in Chapter 4. A comparison of comments across all three tool groups provided support that Athens is more effective at discovering surprising and novel knowledge and reframing mental models than the other tools. Despite this support for Athens, there are still a number of improvements that need to be made. Some of the suggestions made were incorporated into the tool and evaluated in Studies 2 and 3 (Chapter 8 and 9 respectively) – the next iterations of the build-evaluate NKD design-theory process.

One of the limitations of the research design is that each participant only used one tool to complete the task and therefore provided a limited basis for comparison. To compensate for this limitation and provide an alternative method of evaluating the three tools, the product of the task

⁹ In addition, there were scale reliability issues with STERM.

– the novel ideas – were evaluated by five independent judges on radicalness and potential benefit criteria. Although no significant differences between tools were found, it is interesting that the ideas from the Athens tool group were rated higher on radicalness criteria on average. Individual creativity may play a large role in the ideas that individuals generate. However, it is plausible that an NKD tool like Athens can help support the discovery of novel knowledge and the creative process, making the process more effective and efficient.

Another limitation of this study is the use of student subjects. Students, although experienced with the topic of MP3 technology and with Web searching, may not be ideal candidates for completing a level 3 task. It is possible that using organizational participants instead of students may have resulted in a different set of outcomes. However, prior research suggests that the use of student subjects is appropriate (e.g. Greenberg et al., 1987). To compensate for this potential limitation, the case study and field experiment include organizational participants using and evaluating the Athens tool. These results are discussed in Chapters 8 and 9 respectively.

Although Theory 3 results provided support for several of the Theory 1 hypotheses, the results across all three theories indicates that each level of tool can support different types of knowledge discovery and learning. This caveat was noted in Chapter 3. However, each level of tool may be more efficient for certain types of knowledge discovery and learning. For example, Vivisimo's features may make it most efficient for discovering broad knowledge. Vivisimo can help discover deep and novel knowledge, but other tools may be more efficient for these types of knowledge discovery. Technological frames may dampen these efficiency differences in the short run.

As a result of this study, measurement instruments were developed for knowledge-discovery types, learning modes and learning processes. The exploratory factor analysis provided evidence supporting the convergent and discriminant validity of most of the hypothesized constructs. However, there were problems discriminating the mental model maintenance (MMM)

and mental model tuning (TUN) constructs, which may have limited the findings of this study. Future research should focus on discriminating these two constructs more fully. The exploratory factor analysis also suggested some alternative factor compositions that describe different dimensions of the novelty construct, including surprising relationships, new knowledge and novel insights. Future research should develop these dimensions and measurement instrument for the novel construct more fully.

7.6 Summary

This chapter reported the results of a lab experiment in which the effectiveness of three different tools used to complete a level 3 task were compared. Competing hypotheses associated with the knowledge-discovery theoretical framework and technological frames were evaluated. As part of the iterative build-evaluate design-theory process, the NKD design theory and updated version of Athens were assessed. Although positive results were found regarding the ability of the NKD tool Athens to discover surprisingly novel knowledge and reframe thinking, additional improvements to the tool were suggested.

Chapter 8 discusses the results of a case study in which the importance and challenges associated with novel-knowledge discovery are explored and the NKD design theory and Athens tool further evaluated.

CHAPTER 8

STUDY 2 – CASE STUDY

8.1 Purpose

The purpose of this study was to explore how an organization uses an NKD tool in practice. Novel-knowledge discovery was examined at the individual, group and organizational levels. At the organizational level, the goal was to understand why organizations value novel knowledge, how they currently discover novel knowledge and the challenges they face. At the group level, the group processes associated with discovering and evaluating novel knowledge were examined. Finally, at the individual level the way in which individuals currently use tools to discover novel knowledge and, specifically, their experiences with using the updated version of Athens – an instantiation of the NKD design theory – were explored.

In the following sections, the research design utilized for this study is described, including pretest activities. Next, an overview of the case study and data collection is provided, followed by a description of the analysis procedures. The results of this case study are discussed, as well as the implications and key findings for novel-knowledge discovery.

8.2 Research Design

Given its non-commercial status and early state of development, only one organization has been experimenting with Athens. As a result, a single, revelatory case-study design was utilized (Yin, 1994). The exploratory case study consisted of three main components: group interviews, individual interviews and individual tool-use observation sessions. A detailed interview guide was developed for all three components (see Appendix P).

Prior to entering the organization, the interview guide and questions were reviewed by three MIS professors. The guide was updated accordingly.

Multiple sources of data were collected: group and individual interviews with multiple organizational participants, observation of tool use, and review of search results from the NKD tool. A detailed description of the case-study visit, data-analysis procedures and results are provided in the sections below.

8.3 Case-Study Visit

A large pharmaceutical company was the organizational site for this case study. To protect confidentiality, specific details regarding this organization cannot be disclosed¹. Individuals from several different areas within the research and development division of the organization participated in the case study. Two technology-related departments participated: one that investigates emerging technologies and tools, and another department that supports the technologies and tools currently used in the research and development division. In addition, individuals representing different subject-matter focused departments within the research and development division participated. Thus, a wide range of perspectives and roles were represented. Table 8.1 provides a summary of the departments and roles represented in the case study.

Table 8.1 – Summary of Participants

Number of Participants	Department	Roles Represented
2	Emerging Technologies (for R&D use)	Managers (at different levels)
2	Current Technologies (for R&D use)	Managers (at different levels)
3	Various Subject-Matter Focused R&D Departments	Analysts/Searchers

Prior to the site visit, several planning meetings were held with the organizational sponsor. The goal of these sessions was to refine the schedule of events and list of participants, as well as discuss how to focus and prepare for the NKD tool-use sessions. During these sessions a decision was made with the organizational sponsor to focus the Athens queries on the Pubmed²

¹ A confidentiality agreement was signed prior to entering the organization, with strict guidelines regarding disclosure.

² Online database providing access to indexed citations and abstracts of biomedical journal articles (www.ncbi.nlm.nih.gov/PubMed/).

online database, rather than the entire internet. The rationale behind the sponsor's decision was comparability of results. Restricting the queries to Pubmed would provide participants with results directly comparable to those provided in typical Pubmed searches, used frequently in the organization. This would allow participants to assess how novel and useful the Athens results were, compared to current tools.

Two weeks prior to the visit, the researcher sent an email to all participants asking them to think of one or two topics where novel-knowledge discovery was needed in the organization and provide corresponding descriptive terms to be run in Athens. The email also provided instructions for choosing appropriate queries for the tool. All participants provided at least two different sets of descriptive terms (topic areas).

Prior to the visit, the researcher ran these queries on the updated version of Athens. Queries were run with the Pubmed restriction³. In some cases, queries were also run without any restrictions⁴ in order to compare the results and utility of Athens in both scenarios – with and without the Pubmed restriction. In addition, one query was run using the restriction of a competitor's Web site in order to assess the utility of Athens as a competitive-intelligence tool. The results of these queries were used during the tool-use observation sessions.

The organization site visit spanned three days (see Table 8.2). On the first day, a group session was conducted to provide introductory training on the NKD tool and discuss the goals and objectives of the visit. Individual interviews and tool-use observation sessions were conducted on days one and two. Interview questions were open-ended and exploratory. For some questions, participants were asked to describe specific incidents, for example “please describe a time when you discovered something that you considered novel”. During the tool-use observation sessions, individuals were asked to talk aloud to express their thoughts about the use of the tool and the

³ Athens can restrict searches to a particular domain or Web site.

⁴ Search spanning all Web sites on the internet.

results provided. These sessions were highly interactive – participants not only expressed their thoughts, but asked the researcher questions regarding use of the tool.

At the request of the organizational sponsor, an Athens seminar was provided for interested individuals across departments. The session was valuable for understanding the importance of novel knowledge in the organization and potential applications of the NKD tool. On the last day of the visit, a group session was held during which the results of the tool-use sessions and interviews were discussed with all participants. Each participant shared their opinions and insights regarding the NKD tool and novel-knowledge discovery. The researcher also shared what she learned during the sessions in order to validate her thinking with the participants. The next steps for the study were discussed during both the group session and the meeting with the organizational sponsor immediately afterwards.

Table 8.2 – Summary: Case-Study Visit Schedule

Day	Activities
Day 1	Group Meeting/Interview (introduction to NKD and tool, meeting with all participants) Individual Interviews and Tool-Use Observation Sessions
Day 2	Athens Seminar (for interested individuals across departments) Individual Interviews and Tool-Use Observation Sessions
Day 3	Group Meeting/Interview (wrap-up with all participants) Meeting with Organizational Sponsor (to discuss next steps)

Based on the findings and feedback from participants during the visit, modifications were made to the NKD tool-use process and the participants’ queries rerun four weeks after the visit. These new results were shared with the participants and discussed during the final wrap-up session conducted two months after the initial visit. In parallel, transcripts from the interviews – both group and individual interviews – and tool-use observation sessions were created. Each participant reviewed and verified the transcript from his or her sessions.

8.4 Analysis

The analysis process used followed the Miles and Huberman (1994) qualitative analysis guidelines. The process included data reduction (coding), data-display creation to help identify

and validate patterns and themes, and final conclusions. As recommended by Miles and Huberman (1994), data collection and analysis were intertwined. A field report was created during the case-study visit after the individual interviews and tool-use observation sessions were completed. The report was reviewed with participants during the group meeting on the final day of the visit. This review stimulated further discussion and insights from the participants and the researcher. In addition, findings from the initial visit spurred additional refinement of the Athens tool and tool-use process. Queries were rerun and the results were reviewed with participants in a final group session⁵.

Data reduction involved analyzing and coding the interview and tool-use observation transcripts. An initial set of hierarchical codes were developed based on the interview guide, field report, literature review, theoretical framework and Study 1 results⁶. Two different sets of codes were developed; one for the interviews and one for the tool-use observation sessions. After creating the initial set of codes, each transcript was reviewed and codes assigned to text segments. Emergent codes were noted. Transcripts were then imported into NVIVO, a qualitative analysis tool, reviewed and coded a second time using the tool's coding features. During this second coding pass, some minor refinements were made to the codes assigned and additional new emergent codes were created. The second coding pass provided an opportunity to verify the original coding and assess intra-coder agreement (Miles and Huberman, 1994). Appendix Q provides a list of the resulting codes.

After completing the coding of individual and group transcripts, data displays were created for the purpose of assessing trends, patterns and relationships across participants and sessions. Data displays highlighted themes and illustrative quotes across participants. In some cases, the frequency of a theme was also included in the display. In addition, the themes identified in the Athens tool-use sessions were triangulated with the results from the lab

⁵ The final group session did not contain all of the initial participants.

⁶ Athens issues highlighted in the posttest survey of the lab experiment.

experiment – responses to the questions regarding Athens challenges, features and suggestions for improvement. Conclusions were developed based on the patterns, trends and relationships identified.

8.5 Results

The group and individual interview transcripts were analyzed separately from the tool-use observation transcripts. In the following sections, the results from the interviews are discussed, followed by a discussion of the tool-use observation results.

8.5.1 Knowledge Challenges

Participants discussed the challenges associated with novel-knowledge discovery in the organization. Discovering novel knowledge and making novel connections between subject areas was noted as a lengthy and time-consuming process, which often required serendipity. The overwhelming number and breadth of sources involved in the search process results in information overload. Participants also found it difficult to search for novel knowledge because of the uncertainty involved. It is challenging to search for what “you don’t know you don’t know”. Another challenge related to this uncertainty is that often important absolute novel knowledge is protected and not publicly available. The absence of information could be a signal that there is an important opportunity for innovation, or it could mean that someone is already working on an innovation, but keeping it secret. Thus, finding a way to more efficiently and effectively discover novel knowledge was viewed as very valuable.

Participants also described the difficulty of switching from their familiar search behaviours – looking for something specific – to less-familiar behaviours – divergent thinking and making connections between indirectly related areas. One participant commented that it’s more difficult for people to get “input from different unrelated inputs and then making a story of

it, getting signals and recognizing the signals”. Thus, recognizing signals and knowledge as novel and interesting was challenging for participants.

The organization has access to at least 60 different externally-oriented knowledge products, which are used for searching and knowledge discovery. However, these tools are not fully leveraged. Ensuring these sources align to the needs of the business and add value is challenging. Most individuals were only familiar with the tools they used most often and stuck with those, rather than trying new tools and sources. However, less mainstream tools can be a source of innovative, novel knowledge. Thus, technological frames appeared to affect tool use at this organization, and this was viewed as having a constraining effect on novel-knowledge discovery.

In addition to these external knowledge sources, participants felt that internal sources were valuable for novel-knowledge discovery. However, the organization found it difficult to tap into these sources to discover novel knowledge. An example was given regarding how two different areas of the organization, working on seemingly unrelated subject areas, discovered their work was indirectly connected in important ways. Once the indirect connection was identified, novel opportunities were discovered and developed further. However, it was difficult to make these connections. Thus, novel-knowledge discovery was a challenge for this organization – both internally and externally.

8.5.2 Novel Knowledge

Participants discussed different dimensions of novel knowledge, including the importance, when and why it is important, and types of novel knowledge. Each dimension is discussed below.

8.5.2.1 Importance

All participants agreed that novel knowledge was critical for the organization. In the research and development division, novel knowledge was crucial for innovating – creating a new market, product or service – and viewed as a necessity for survival. The importance of innovation was further highlighted by the fact that it has been incorporated into personal-development plans for research and development employees. In other words, individuals are evaluated on their contributions to innovation in the organization.

Another interesting observation by a participant was that it is often important to find out what you thought was novel is in fact not novel, and vice versa. This has implications for patents and other intellectual-property issues, as well as determining which opportunities to pursue.

8.5.2.2 When and Why

Participants were asked to comment on when and why novel-knowledge discovery was particularly important. Novel-knowledge discovery was viewed as being most important early in the research and development process and during mid-development. In the early stages of research and development, novel-knowledge discovery is important for identifying potentially-valuable opportunities to pursue. In the mid-stages of development, novel-knowledge discovery was deemed important prior to the clinical trials stage – the beta-test of the product. Discovering possible, yet as-of-yet unknown, signals regarding safety issues before launching a trial was viewed as very valuable since the cost of making an error and “scrapping” the project at this point is extremely high.

On an ongoing basis, novel-knowledge discovery was viewed as important for recognizing unknown connections between internal knowledge areas, for keeping up-to-date with emerging technologies and science, and for competitive intelligence.

8.5.2.3 Types

One of the emergent themes uncovered during data analysis was the different types of novel knowledge relevant to the organization. In addition to the types of novelty identified in the literature – relative and absolute novelty – several subtypes of novel knowledge were identified: novel connections or associations, new root causes or new solutions to old problems, new problems, and new applications for an existing product or service.

8.5.3 Novel-Knowledge Discovery

Several broad themes related to novel-knowledge discovery were found in the interview data related to processes, types of searching behaviour, outcomes and sources. Each of these areas is discussed in further detail below.

8.5.3.1 Processes

Participants were asked questions regarding the process of novel-knowledge discovery and asked to provide examples of cases where they discovered something they perceived as novel. Several process-related themes were discovered across the responses.

An interesting finding regarding the novel-knowledge discovery process was that there really wasn't a distinct process that individuals followed. Individuals don't tend to initiate a search with the specific goal of discovering novel knowledge. Two individuals commented,

“I don't think people think 'I'm going to go do novel discovery'. They don't think of that as a distinct process. It's just keeping up with... That's one of the things on my development plan - things I have to do to keep up with what's going on in the industry.”

“There is no focused effort just to do novel knowledge discovery. I think virtually everybody is really interested in that kind of discovery.”

However, when individuals do discover something novel, they describe it as being the outcome of the following activities: scanning the literature, using alert software to keep up-to-date on the latest developments, searching for something specific and stumbling upon something novel, and

networking and discussions with colleagues – both internal and external to the organization. One participant commented,

“Because I was looking for something that I had to look for in order to do my job. I wasn't specifically going out to find new knowledge. I was going out specifically to address a very specific problem. And then secondarily that new knowledge came about because I was reading more broadly than I had to in order to get the answer.”

Thus, a variety of processes with different initial motivations – goal-based, problem-based or keeping up-to-date – led to novel-knowledge discovery. Serendipity is almost always involved in the discovery of novel knowledge.

Two other key themes about the overall process of novel-knowledge discovery involve groups and change management. There are groups within the organization tasked with the role of keeping up with the latest developments in an area. These groups may be important and fertile grounds for discovering novel knowledge. Both groups and social networks were noted as being important in the process of generating, evaluating, integrating and institutionalizing ideas. However, groups can impede the institutionalization of novel knowledge when knowledge ownership issues are salient.

Change management was noted as another important component of novel-knowledge discovery. People are generally resistant to change. Several participants discussed examples of how particular novel ideas were successfully implemented in the organization. In many cases, the challenge regarding adoption was the sheer novelty of the idea. In order to overcome resistance, reframing was used as a technique to allow participants to understand how to approach and use the new process, product or technology. In other words, the introduction and training provided helped reframe mental models of how things are currently done so that individuals could envision how the new process, product or technology could be successfully integrated and used in their work and across the organization. Thus, the introduction of novel ideas was challenging and required proper management and framing of the change.

Participants also discussed the use and importance of tools in the discovery of novel knowledge. In general, tools were viewed as critical and a key resource for novel-knowledge discovery. The following comments illustrate the importance of tools:

“So a lot of these technologies really do have an impact on and the overall big picture of how we can be innovative.”

“Well we are surrounded by tools and it's a matter of knowing first of all that they exist and then knowing how to use them. So yes tools. I try to use as many tools as possible.”

Participants also noted the importance of social networking and discussions with colleagues as part of the process of novel-knowledge discovery. Thus, tools are deemed an important resource for novel-knowledge discovery, but social networking is an important component as well.

Interestingly, there are no specific tools used for novel-knowledge discovery at the organization. The organization has a large number of knowledge sources and tools. Novel knowledge comes from all systems, usually through serendipity. The types of tools utilized include alerts, search engines such as Google, access to external databases and journal articles, internal document repositories, internal thesaurus, internal document-indexing and search tools, and collaborative software such as social-network-analysis tools.

Participants discussed the process they typically followed when they found potentially novel and interesting knowledge. Initially, raw novel findings⁷ may be shared with others – either with an individual or a group. These results may be shared either informally or formally, depending on the context. If warranted, results are translated into an idea – either initially, or after sharing the results with others. The translation of the novel findings to an idea involved taking the findings and identifying an application viable for the business – either the existing “as-is” idea or an adaptation for other purposes. This process of translating or adapting the novel findings to something that can be used by the organization involves thinking about the idea, articulating the idea and “getting a good grasp on it”. Once an idea is generated, it is discussed more fully with

⁷ For example, a link to a Web page or a journal article.

others, analyzed and evaluated. The idea may be shared with individuals or groups to whom the idea may be valuable. Formal evaluation of the idea involves a cost/benefit analysis. The degree to which the idea is novel is never a criterion in this analysis. However, the novelty may result in a tangible or intangible benefit. The feasibility of the idea is also considered. Formal evaluation can be an individual or group process, depending on the context.

8.5.3.2 Types of Searching

Analysis of the interview data revealed several different types of searching behaviour, further distinguishing focused search and scanning into finer levels of search. In terms of focused search, the following different subtypes were identified: retrospective versus current search, focused and comprehensive search, focused search for a specific answer, and focused search for competitive intelligence. In addition to general scanning behaviour, the subtype of monitoring with alerts was identified. This subtype of scanning can be further broken down by the breadth of subject – general or specific. Table 8.3 provides a summary of these search subtypes as well as illustrative quotes. Discovering novel knowledge is a by-product of these searching behaviours. There is typically no explicit intention to discover novelty.

Table 8.3 – Summary of Search Subtypes

Search Type	Search Subtype	Illustrative Quotes
Focused search	Retrospective vs. current search	“The sort of work that I do, which is mostly retrospective and highly comprehensive in nature.”
	Focused and comprehensive search	“The kind of searching that I do, for the most part, is extremely comprehensive. You’re not just looking for a few good references. You’re looking for anything and everything.” “I don’t use Google for my comprehensive searching for the simple fact that you don’t always get the same thing twice. And I need to get consistent results. But there are times when it takes you right to what you need to know to supplement. I use it for additional information rather than the main thing. Yes it can be useful.” “It might not be important for something like this, but when you’re doing comprehensive searching it’s important. The trouble with this type of searching in real life is that if you put <search term A> and <search term B> too close together you will miss stuff and if you put them too far apart you get <search term A> for other stuff. It’s very prone to false drops. Very voluminous searching to do.”
	Focused search for a specific answer	“And you start off with just a simple search and if you get the results you want great. But most often we are asking questions

Search Type	Search Subtype	Illustrative Quotes
		as precisely as possible. So we only get back what we want." "I guess I generally do not look broadly. I am usually so focused on identifying a solution to a particular problem. And then new knowledge comes along." "Typically when I am going out to the web and looking for something I am looking for something specific. So if I find an exact hit I will copy it and paste it and essentially I'm done."
	Focused search for competitive intelligence	"I also sometimes try to go on the Internet to check company Web pages. I use the Internet in this case. I run a search using Google. But not the other one that you mentioned this morning, Vivisimo. It's the first time that I have heard of this name. I of course know Google. I don't run searches here generally. I just use it to look for companies and the web pages of companies." "Those requests can be...competitive intelligence on competitor assets or competitor companies."
Scanning	Monitoring with alerts	"One of the things that I do is that I monitor the news. We have this system called Factiva, which is like a Dow Jones. I have some alerting profiles set up in Factiva where I am watching for events related to patent disputes. Sometimes it can be quite specific where I am looking for something that happens with a particular product - our product or somebody else's product. Or sometimes it can be very general in terms of who is suing whom, is there some change in legislation that we're interested in, or a dispute of particular interest." "The other thing is trying to be more proactive and trying to stay up to date. So setting up alerts." "RSS aggregators are usually thought of for news or for blogs, but it's an ideal tool for delivering published information alerts. People set up alerts for things like industry newsletters..."

The analysis also revealed that the search process itself is iterative and emergent. The process involves going down different paths that seem appealing or interesting. Individuals assess both the relevancy and interest in the results. The initial assessment can involve a quick scan of the results. If the individual finds something interesting or equivocal, she or he delves down into the results in more detail. Search terms can be further manipulated in order to focus and filter the results. Although the context of the search affects the exact process followed, these components were the generic search and analysis processes discussed by participants.

8.5.3.3 Outcomes

During the group interview on the final day of the case-study visit, participants discussed both the types of knowledge produced by Athens as well as the tool's usefulness. Participants found that the tool produced both incremental knowledge and novel knowledge – both of which are needed for organizational adaptation (Edmondson, 2002). The novelty of the results

depended on the individual's current knowledge, amongst other things. Individuals noted that the tool identified some interesting results, which could generate novel insights. However, they also noted that the tool produced a lot of uninteresting and irrelevant results as well.

Given the Athens issues identified in the tool-use observation sessions (discussed further in the next section), the Athens results reviewed during the case-study visit were not viewed as particularly useful or valuable. Despite the problems with the NKD tool "in practice", participants felt that the concept of what Athens was attempting to do – "in theory" – was very useful and interesting. Several participants were eager to try the tool again after refinements were made. Participants also suggested ways in which the tool could be applied in the organization. For example, one suggested application of the tool was to link it to the electronic alerts systems so that individuals could be periodically alerted when their NKD tool produced new novel results.

8.5.3.4 Sources

During the initial group meeting, there was a lot of discussion about the sources that could or should be used for novel-knowledge discovery. Several different types of sources were discussed: literature databases, journals, patent databases, the Web in general, competitor Web pages, news Web pages, blogs and internal data sources. In addition to the type of source, both the age and structure of the source were discussed in terms of the ideal fit for novel-knowledge discovery. Sources differ by the age of content. For example, online journal publications may be a year or more old due to the publication lifecycle, whereas blogs and news sources are very recent and up-to-date. However, both older and more recent sources could be useful for novel-knowledge discovery, depending on the context. Sources also differed in the degree to which the content is structured. Literature databases such as Pubmed are highly structured sources, with extensive indexing functions and very limited free text. Other sources, such as the Web and journal publications, are much less structured with limited indexing, yet extensive full text content. Using Pubmed for novel-knowledge discovery in the tool-use observation sessions

proved to be somewhat challenging (discussed in more detail below) due to the limited amount of full text available. Thus, less structured sources with more full text may be better sources for novel-knowledge discovery.

8.5.4 Innovation Approaches

During individual interviews, participants discussed the various approaches to innovation employed by the organization. The different types discussed were sole-development, co-development and purchase of licensing rights. Finding suitable partners for co-development is important for innovation at the organization. This is an area where a novel-knowledge discovery tool could add value by identifying interesting and novel partnering opportunities.

8.5.5 Athens Tool-Use

Analysis of the transcripts from the Athens tool-use sessions led to the identification of the following high-level themes: types of problems encountered, types of learning processes, and overall positive tool experiences. The codes developed for the types of problems encountered were derived from the qualitative results from the lab experiment. Thus, the results described below help support and explain the Athens issues identified in Study 1.

8.5.5.1 Positive Tool Experiences

Despite the challenges encountered while using the Athens tool, there were a number of instances where participants had positive experiences and comments regarding the tool. The themes identified across these positive experiences included finding a result that is interesting and making sense of a novel finding. Several participants commented,

“Yes, that one actually works for me. That is relatable back to what you were... it is off on a tangent, but there is a connection there.”

“Some of it looks like it might be interesting though. Alright, now I could see how this could get this kind of thing. This is an algorithm that is related to networks it looks like. But that is something that could turn out to be related to medical informatics. It’s coming out of the computational geometry area. But I could see

how that might actually be ‘huh, I wonder how that could be usable in some of the medical informatics where we're doing things like laying out results for people’.”

“I think this one is actually the most interesting - the one against Medline with the medical informatics. Is the one that was ‘What are they doing in Japan?’ is what it led me to. It would be interesting to go... that there may be a research group there doing a bunch of interesting things.”

For these particular results, participants were able to make sense of the novel connections and could see how the result could be useful in the context of their topic area.

Other themes related to positive tool experiences included comments about features of the tool participants particularly liked – for example, tag clouds. In addition, participants commented that they liked the overall idea of Athens and what the tool is trying to do. Illustrative quotes include the following:

“But conceptually I like the idea of being able to plug in two or three concepts and broadening those out one or two levels and looking at the intersects. Because generally that is where new knowledge tends to come is the intersects of what has previously been a disconnected set of concepts. And the two come together and all of a sudden somebody has an idea, a new technology, or new product.”

“I am actually fairly eager to try it again <without the restriction to Pubmed> and more of a chance to try some different queries. I can see that it might be potentially very useful in the areas we work in. For some relationships that we might not have thought about between new technologies and what is going on in different domains.”

“I think it was interesting, the whole concept and the whole process.”

Participants also commented on the challenges and issues they encountered while using the Athens tool. These issues are discussed in detail in the following section.

8.5.5.2 Types of Problems

Many of the same problems discussed by the participants in the lab experiment were identified while analyzing the Athens tool-use transcripts. In Appendix R, details regarding each problem found in both the lab experiment and the current study are provided, including both the number of comments and participants who made these comments. The problems not mentioned in this study, but mentioned in the lab experiment and vice versa are also noted.

Similar to the lab experiment, participants found it difficult to navigate through the Athens tool results. Part of this was due to the novelty of the tool itself. Participants made several suggestions regarding improving the usability of the tool, such as including a more visual display of the results. Additional visualization and meta-data were proposed as mechanisms to help improve the user's ability to scan and interpret the results more quickly.

The results from the tool-use observation sessions surfaced an important issue, which impacted the usefulness of the results. During these sessions, it was discovered by several of the participants in collaboration with the researcher that the tool was picking up proper nouns, such as the names of people and places. As part of its processing, the Athens tool uses text-mining techniques to identify the most important nouns in the documents retrieved. Part of the logic used to assess importance is the comparison of the prevalence of the word in everyday language to the prevalence of the word in the retrieved documents. Since nouns such as the names of people and places are relatively less common in everyday language, the tool was identifying them as important in the retrieved documents. This issue was more prevalent in the results that were run without the Pubmed restriction. For example, the tool would flag the noun "Florida" as important. Thus, some of the novel results were skewed to topics related to Florida. These particular results appeared to be completely irrelevant. Another related problem was detected in the Pubmed-restricted queries. As participants reviewed the Pubmed-restricted results, it became apparent that the nouns describing each cluster were almost identical across all clusters – for example, gene, genome, protein, RefSeq, books, Genbank. Prevalent Pubmed nouns, which are very uncommon in everyday language, were being highlighted as important. Thus, some of the results were being skewed by these common Pubmed nouns and, thus, less interpretable.

The issue with proper nouns was also indirectly mentioned by a couple of the participants in the lab experiment – for example, a comment from one participant in the lab experiment that the results were in German. However, the level of interaction during the tool-use observation

sessions in this study helped the researcher better understand the issue and initiate a search for solutions.

Another dominant theme in the results of the tool-use observation sessions was the existence and influence of technological frames. Comments about the Athens tool appeared to be strongly influenced by each participant's familiarity and experience with other search tools. For example, several comparisons were made to Google. In addition, some participants were clearly trying to use the Athens tool as a focused search tool. Illustrative comments follow:

"I think because the smallest thing takes 45 minutes to run and it can take up to ten hours... mostly people are looking for instant gratification here and clearly you would have to think out your query carefully well in advance and be prepared to wait that long. It's not like Google where you Google something and it flies up on the screen."

"Yes. I am curious to see what this did because I also picked this for a reason. Because it gives a lot of false starts...If you use of a big boolean AND, and you say give me everything with <search term A> AND <search term B>, you're going to get <search term A> and everything else under the sun but you don't want. I didn't know whether Athens had the capability or not to figure that out."

Some participants had clear expectations about what they were looking for and what they wanted to find. In fact, a common remark was "Yes, that is interesting, but not for what I want or what I am looking for". Thus, similar to the lab experiment, some individuals were using Athens for a focused-search task rather than using the tool in an exploratory manner where surprises are desirable. One of the participants actually recognized that individuals were using the tool with focused-search based expectations.

Other concerns noted by the participants include irrelevant results, errors, search-term refinement, quality of the Web links, and difficulty understanding the connections between the clusters and the original search terms. Some of these problems were the result of the proper noun issue and some the result of the prototypical state of the tool (for example, errors and the need for further development of the user interface). More light is shed on the concerns regarding understandability of results in the following section, in which learning processes are discussed.

8.5.6 Learning Processes

Themes related to learning processes were identified in both the interview data and the tool-use observation data. The initial codes created were based on the theoretical framework and updated 4I organizational learning process model. However, two emergent themes were discovered during the analysis: goals and context. It was evident in the transcripts that the search and learning process started with a goal. How the individual searched and what the individual learned from the results depended on the context of the search task.

The definition of search terms is based on the goals of the search. However, as identified in the lab experiment, defining search terms is challenging when you do not know what you are looking for – a challenge endemic to novel-knowledge discovery. Defining and refining search terms is also difficult when the individual is not familiar with the subject area. The importance of individual knowledge and expertise was evident from the analysis of the interview and tool-use data. Some of the participants were not completely familiar with the subject matter of their search, which influenced their ability to choose search terms and interpret the results. In some cases, individuals could not comment on the novelty of the results because they weren't experts in the topic area. In other cases, participants commented on the difficulty of selecting search terms because of their lack of familiarity with topics appropriate for searching within Pubmed. One participant commented,

“It's interesting in that it's pointing out that if I hadn't gone in knowing already that a lot of what computing is doing right now within health care has to do with genome and protein sequencing. That is what is popping out here - genes and genomes. And this is taking me in a direction of...”

Thus, this participant would not have viewed this particular novel finding as relevant without a certain level of expertise in the area.

Additional insights were gained into the process of information foraging. In general, information foraging is an iterative process that involves collecting a large amount of data and

information, determining which paths to follow and which places to go into more detail. Most of the participants were experienced searchers and understood the science of searching. When analyzing search results, individuals reported that they tend to either review in depth or scan broadly the top results. Individuals assess the relevance and decide whether to stop or continue searching through the results. Participants indicated that they usually stop if they find the answer they were looking for. However in some cases, the search is continued if the individual continues to find interesting things. One of the participants commented,

“If I can find the answer then obviously I am going to stop there. Because generally I am looking for a specific topic. Probably 50 percent of the time I will look at other pages. So I might go to the next page... so page two, maybe page three, and if I am still finding stuff that is remotely interesting then I may jump down several pages... so page five, page ten, just to see what is available there.”

“Sometimes you will hit a gold mine and say ‘I don't need all of this stuff right now. I only need 1 % of it. But it is a resource I can go to in the future’. And because of the Web, these things tend to reference other sources. So you tend to spider out and kind of the mushrooming effect until you find what you need.”

There is often a lot of noise in the information retrieved, for example, irrelevant results. Thus, the information foraging process is iterative and involves manipulating the search terms, finding synonyms and filtering the results. Individuals will adjust the search terms if they discover they are not on the right path,

“I will, depending on the time that I have available, I will bounce around further down the list. And if I realize that I am not on target, I will reevaluate the way I did the search. And try to be more broad, or more precise. And adjust accordingly.”

If individuals perceive they are on the “right” path, they may simply adjust the search terms to “prune”—refine and filter – the results,

“And you ultimately get through that process of searching and looking and fiddling around more and getting it to a reasonably comprehensive yet not overwhelming result set. And then exporting it and maybe having to edit out any other outliers that couldn't be reasonably gotten rid of through the strategy itself.”

In most search tools, this adjustment process is manually performed by the individual. However, with Athens most of this process is performed by the tool, which causes its own set of problems

regarding use of the tool. People are used to having this level of control with focused-search tasks. The attempt by some participants to use Athens as a focused-search tool highlights this desire for control of the search-term refinement process.

The tool-use observation transcripts provided rich data for exploring and understanding the intuiting and interpreting processes of learning. Intuiting and interpreting processes were difficult to differentiate from one another. Defining the constructs ahead of time and using coding rules to distinguish the two processes was an important part of the analysis process. Carefully analyzing equivocal text segments helped refine and augment the definitions, resulting in construct definitions grounded in both the literature and data.

Intuiting, as discussed in Chapter 2, involves pattern recognition and the development of insights into past patterns and future possibilities. Ideas are not yet formed, but are subtle, fuzzy, initial inklings of interesting possibilities (Crossan et al., 1999). The fuzzy and subconscious nature of intuiting made it challenging to code and analyze the transcript data for intuiting learning processes. However, this definition was used to help identify and code intuiting processes in the data. A number of intuiting-related patterns were found during the analysis, helping to elaborate the definition of intuiting and the challenges associated with intuiting and novel-knowledge discovery (see Table 8.4).

Table 8.4 – Summary of Intuiting Themes

Theme	Illustrative Quotes
Thinking about patterns, signals	<p>"But there are some interesting things in here also. Like wiki, Disneyland... I think what we are clustering here is geographic."</p> <p>"So that we can either research them or just scan over them. Sometimes by just scanning over them you will have additional questions that might spark additional searches or the generation of new knowledge."</p> <p>"Okay. So some of these are obviously well connected to the search terms. Others are already a little bit remote as far as the two concepts."</p>
Initial thoughts regarding interesting possibilities	<p>"There's clearly something going on in Japan that's of interest."</p> <p>"My customers, they think that maybe in the early stages that there is a target use for a specific indication."</p>
Considering, questioning	<p>"Genes, no. Why is there so much information on genes?"</p> <p>"Real estate?"</p>
Generating an understanding of why	<p>"Looks like the focus here is primary care."</p>

Theme	Illustrative Quotes
patterns are emerging	<p>“This looks like what it is getting at is primary care or home care kinds of things.”</p> <p>“The places are probably either conferences or universities.”</p> <p>“Some of it looks like it might be interesting though. Alright, now I could see how this could get this kind of thing. This is an algorithm that is related to networks it looks like. But that is something that could turn out to be related to medical informatics. It’s coming out of the computational geometry area. But I could see how that might actually be ‘huh, I wonder how that could be usable in some of the medical informatics where we’re doing things like laying out results for people’. I don’t know why he would be brought in. And that one I think is not particularly... this is getting into science fiction.”</p> <p>“Yep. Yes that is definitely relevant. I wonder if it is just Science that is doing this to it. Maybe there are some issues with Science. There, again. Is there any importance to the order in which these are listed?”</p>
Deciding where to go next	<p>“When you get the results from your query, and it sounds like it’s very iterative, you get some results in and say ‘hmmm I want to filter it more or I didn’t get what I wanted’?</p> <p>Yes. Sometimes the stuff is wrong somehow, or there’s too little, or there’s too much. I fiddle with it a bit. And what often happens is I will go and read the original documents.”</p> <p>“It might be interesting to look at what’s going on in between.”</p> <p>“Interesting. Computer center. Looks like an interesting place. Here we go. That is interesting actually because it’s leading me to look at some of the other things they are doing there. Yeah. I will look at this more later. So even with the places and things in, it has led me in some interesting directions. Like they are doing a bunch of interesting work there.”</p>
Reducing uncertainty – looking at more details to get a better understanding of the pattern	<p>“Actually this one time we were researching something on a regulatory action that just didn’t make sense and nothing in my regular sources was pulling up anything that explained it. I couldn’t figure out why they’re getting this extra regulatory protection. It didn’t make sense... So I decided to try Google. And he said ‘you’re Googling it?’ And I said ‘why not nothing else is working’. And we found out exactly what was happening.”</p> <p>“It’s generally a quick scan of the title to start with and then maybe the abstract and any kind of key words.”</p>
Encountering, triggering or thinking of new ways to attack the problem	<p>“But it’s not at all unusual, particularly if it’s a new subject area, that I start going through these and find additional synonyms that didn’t occur to me the first time. Other names for these substances, other names for this disease condition, that I hadn’t thought of.”</p> <p>“So pull them down in different ways. So maybe they do a search and that triggers something - I should look here and go talk to this person. So it’s a combination of a lot of different things and some serendipity thrown in there as well.</p> <p>Yes.”</p> <p>“Sometimes it’s frustrating when you don’t find the answer. But sometimes it forces you to generate the answer yourself and generate a little additional knowledge at the same time.”</p>

These intuiting themes are related to initial thoughts about points of interest, directions and uncertainty. Individuals are trying to clarify, understand, and reduce the uncertainty with this

initial thinking. As many of the quotes above suggest, intuiting is related to and interwoven with the information-foraging process. Intuiting is required to make decisions about which paths to follow and where to look for more details. Ideas are not yet formulated in the intuiting processes. The novel insights developed in the intuiting process are “crystallized”, developed and articulated as ideas during the interpreting process (Bontis et al., 2002; Ghiselin, 1952).

In addition to the formulation of ideas, the interpreting process involves developing models for understanding – new or updated mental models – and bringing out meaning. The generation of ideas was not the goal of the tool-use sessions, thus most of the learning processes identified were related to intuiting. However, there were three themes identified as interpreting processes (see Table 8.5).

Table 8.5 – Summary of Interpreting Themes

Theme	Illustrative Quotes
Generating/translating an idea	<p>“So maybe they will start thinking about developing the proposal for a new drug.”</p> <p>“No. It’s been in the news. But it may tell us something when ours is approved, what approach we might... may tell the people marketing and promoting this.”</p> <p>“It’s not new information. But as a result of this, we may moderate are own, tailor our own approach to marketing this.”</p> <p>“I like to get a good grasp on an idea before I go ahead and release it. And that is where you have to do a little bit of work to translate that similar idea to meet your needs. But by having something that is similar, it means that it is not exactly what you want. So in general you are going to learn something else from that exercise.”</p>
Making an assessment or decision	<p>“No. Seems like there is a lot on genetics, but it is not related. Not information that is useful in this case. Maybe in some others, but not this one.”</p> <p>“You are telling me that the part that is related to food and genetics, something novel that can be applied to this area? Correct? This can be very useful information but not in my case for the requests that I got.”</p>
Making sense	<p>“And when I did the search I found a reference from 1964 that described it. Now that’s not new information, but it was new to us. What it told us was that what this company was calling novel was in fact not novel. This is very interesting because it says that the patent that says this is new, they’re not going to issue. They apparently didn’t find this reference before they filed the patent.”</p> <p>“Actually that is kind of interesting as it suggests that AI had something to contribute in the whole area of home care if we looked at it.”</p>

Some of the interpreting themes identified above are highly related to the intuiting process themes. It is arguable whether any interpreting processes were involved in the Athens tool-use sessions given the lack of focus on idea generation. It is, therefore, plausible that these interpreting themes are actually additional intuiting themes, which may lead to interpreting processes eventually. In either case, intuiting and interpreting processes are highly related and frequent iterations between the two processes are likely.

Integration of the learning-process themes with the Athens problem themes provided insights into how tool support for learning processes may be enhanced. Based on the comments of the participants, visualization may help with intuiting processes by making patterns easier to see and making connections between different areas more apparent. Additional meta-data may also help intuiting processes by better enabling individuals to detect areas of interest and decide where to go into more detail. These suggested tool enhancements may help to focus the user's attention and reduce "noise" in the results. Improving the filtering of results, especially proper nouns, should help interpreting and sense-making processes, leading to the generation of ideas. While it is the individual who ultimately generates the idea, the novel clusters and the indirect connections between subject areas may help spur the generation of novel ideas just as indirect stimulus enhances creativity (Hender et al., 2002; Nagasundaram and Bostrom, 1995; VanGundy, 1988).

Although the group and organizational learning processes were not part of the Athens tool-use sessions, evidence of these processes was found in the individual and group interview data. Three themes were found in the interview data regarding the group-interpretation process. One of the group-interpreting themes involved the generation of a group understanding of an idea. One participant commented,

"And one thing that we had to do in early negotiations with <external software vendor>... on the interface itself they were referring to things that were called posts. In the blog arena that's what they are. But I'm talking about scientific journal articles. Who would ever relate a journal article to a post? And they said no, we cannot do that because you're the only one with journal articles. Well what about news articles and if you want to really get big in the corporate

enterprise market then you might want to consider that. They said they would compromise and change it to items. But I wanted articles...”

“I was just trying to influence them and how they were thinking about the way the content is being delivered through their software. So that was my innovation, but it came out of left wing. People were like ‘RSS for published articles?’ But they got it then.”

A shared interpretation of the idea was not immediate in the situation described above. Debate and discussion were required to develop a shared and negotiated understanding.

A related theme is that of group influence. The idea itself can change as members of the group shape the understanding of the idea. One participant commented,

“And whenever you can actually sit down and talk to somebody in a casual fashion, inevitably you will share ideas that will propagate other ideas. Almost every time I come away with a new idea.”

“But then I would talk about it in a group environment or to other specific individuals who could benefit or give me input on the validity of the idea.”

Thus, the group can influence and shape the idea through group discussion – the final theme related to group interpreting.

Themes regarding group-integrating processes were also found in the interview data. The group integrating process involves taking action on the idea in an effort to implement it. Two integrating themes were found, the first being implementation and adoption of ideas. The following comment is illustrative of the processes and challenges involved with idea implementation and adoption:

“A lot of times when we look at alerting, you look at all these systems eventually coming into the email. We knew it was innovative and thus a lot of change management would need to be involved. That’s a big deal in product management - change agents... The biggest disincentive for using this RSS aggregator is that they are not getting their updates in their e-mail any more; they live and breathe in their e-mail... We had to identify the advantages of using this. You had to go to another system to find this, but there are so many more advantages to doing it that way than having all these disparate emails coming in at different times a day and the week and the month. You are more in control. And once people see it they get it. We actually had this best practice e-mail go out telling people now that you have this you have to shut off your e-mail alerts and here's how. So all these people were shutting them off and they're doing that primarily because they got it.”

Thus, the novelty of an idea affects the ease of adoption, similar to the influence of technological frames discussed in Chapter 3. Managing the change and understanding the obstacles is important.

The other theme related to integrating processes is pilot and evaluation processes. Pilots are useful for evaluating the instantiation of an idea prior to full roll-out and identifying possible obstacles, such as those identified above. The following comments address the use of pilots in the organization:

“What we ended up with was the project where we brought in wiki and we did first just a rough draft with a few people in that group to try it out. Then we did a pilot where we tried it out for several months.”

“It took about a year and a half to get from ‘yes we will look at this idea’ to ‘we’ve handed it off’. Sometimes we look at things for 2, 3 or 5 years and keep trying different pilots and ‘yes this really does look valuable, but we haven’t found the right place in the business for it yet’. Or ‘yes this really looks like it will be valuable eventually, but it isn’t ready to put in a full fledged production system’. So there is variety of different processes for how we take things forward.”

Once an idea is developed and implemented, the process of institutionalizing the idea begins. During the interviews, it was apparent that the organization uses a “grass roots” method of getting the organization to adopt and buy into an idea. The following comments discuss how ideas are adopted and diffused in the organization:

“Within this company we have things like the emerging technologies council... My department...is not really members in that, but we hang out with people like <names removed>. So people we know and we work with that are involved with that can say ‘hey this RSS stuff... we have an idea that we could actually use it for this’. And then we will go make a presentation to the group and then it takes off. In fact since we launched this and people found out that we had an RSS reader, people have been coming to me. The committee is responsible for this. There is a group called the Internet Steering Committee (ISG) and it’s at the corporate level. They got in touch with me and said they were investigating RSS aggregators as a possible supporting technology that can be used for managing news and gadgets; and they will pull it all over to our web pages... Now they are doing a pilot. So now there are representatives from all these different departments and groups globally. They are being exposed to it that way.”

“There are a lot of committees that bring people from different groups together to share ideas and show ‘hey we did it this way and here’s some ideas of how else it

can be exploited. Are you interested, do you want to try it?' And then they start playing with it.”

In addition, ideas are institutionalized by embedding them into routines. One participant made the following comment,

“Most of the people using it now don't even know what RSS is. It's not related.”

This comment illustrates the taken-for-granted nature of institutionalized ideas.

8.5.7 Secondary Athens Evaluation

One of the most significant issues identified during the tool-use sessions was the influence of proper nouns on the results. Participants were intrigued at the potential of Athens, but the influence of proper nouns was making it difficult to get a true evaluation of Athens usefulness. The organizational sponsor made the following comment at the conclusion of the visit:

“I am really kind of intrigued. Because I think in some ways we just didn't conduct the test particularly well. It's fine, this is our first cut. So I think we're all learning a great deal about this. But it would be interesting now... now that we know more, to focus Athens on a better source... and filter out these proper names and a few things. And I think with a few things you could start to see whether or not this system has some nice capabilities.”

The study participants and organizational sponsor suggested conducting a second evaluation of the tool once proper-noun filtering was added.

After completing this initial visit, the researcher determined that in the short run, automating a proper-noun filtering mechanism in the tool was not feasible. Instead, the tool's “stopwords” list was used to manually filter out proper nouns. The stopwords list contains words that are purposefully excluded from Athens analysis of important nouns.

As requested by the organizational sponsor, the Athens queries used in the initial visit were also used for the secondary evaluation in order to provide a basis for comparison. The process used by the researcher to filter out proper nouns involved a two-step process. In the first step, the researcher ran each query in Athens and retrieved the list of “important” nouns that

Athens used during processing. The researcher manually reviewed the list from each query and noted which words were proper nouns that should be added to the stopwords list. Each query had its own unique stopwords list. In the second step, each query was rerun with the updated stopwords list. The new Athens results were reviewed in depth by the researcher and compared to the original Athens results used during the case-study visit. The new and original results were sent to the participants for review.

A meeting to collectively discuss the results of the secondary evaluation was held two months after the date of the initial case-study visit. Comparing the original and new Athens results, the differences and improvements were evident. The participants commented,

“The results are more relevant to the domain.”

“I can see relevance but can’t make the judgment regarding the novelty of the results.”

“I can see logical connections to some of the novel stuff.”

Additional comments were made by the participants regarding how to further improve the results, many related to the themes discussed in the initial visit. Several suggestions were simple ways to make the output more navigable.

During the meeting, the researcher shared her insights regarding the results and the process of manually filtering the proper nouns. The process of choosing stopwords to add to the list – which proper nouns are interesting, and which are not – was a subjective process best performed by the individual. Fully automating this process may result in filtering out some proper nouns that are important in the context of the domain.

The process of selectively adding proper nouns to the stopwords list appeared to have improved the results. A number of different effects on the results were found. In the original results, participants often commented on the results being too novel – “too out there”. Adding stopwords helped reduce the extreme novelty of the results. Results weren’t as “out there” and appeared to be more relevant and connected to the key terms, yet still novel. The nouns used to

describe each cluster were now more meaningful. In Pubmed, the Web links returned tended to contain more meaningful content as opposed to lists of citations and other types of less meaningful document records. In some cases, the Pubmed-restricted results contained more tightly-clustered results, making it harder to distinguish clusters from one another. However, this did not happen in all cases and may be based on the scope of the original search terms. Overall, the effect of adding stopwords improved the interpretability and usefulness of the results.

In general, the participants agreed with the researcher's insights. The organizational sponsor's final comments included "pretty interesting results" and "I'm impressed by the relevance of the results". The sponsor also noted that it was "great that doing something relatively easy, like adding stopwords, could help improve the results". The sponsor expressed interest in pursuing the development and use of the Athens tool further and suggested that the results be shown to the research scientists within their organization – a positive indication of the potential value of an NKD tool like Athens.

8.6 Discussion

The purpose of this case study was to explore the practice of novel-knowledge discovery at the organizational, group and individual levels. At the organizational level, novel-knowledge discovery was viewed as critical to the innovation process. However, there were no distinct processes in place to discover novel knowledge. Rather, a variety of processes with different initial goals led to the serendipitous discovery of novel knowledge. Novel-knowledge discovery was imbedded within and a by-product of these processes. A number of different applications of novel knowledge in the organization were suggested – identification of opportunities and risks, identification of partners, and competitive intelligence. Further, a number of different types of novel knowledge were identified – novel connections or associations, new root causes or new solutions to old problems, new problems, and new applications of an existing product or service.

A number of different tools were used across the organization for searching and knowledge discovery. Aligning these tools to the needs of the organization was a challenge. Despite the importance of novel-knowledge discovery, there were no tools used specifically for this purpose. Any and all tools could lead to the discovery of novel knowledge. However, discovery was unintentional and serendipitous. Technological frames appeared to have a constraining effect on novel-knowledge discovery due to the tendency of individuals within the organization to use only the tools with which they were familiar.

Support was found in this study for the NKD design theory. At the individual level, the challenges associated with novel-knowledge discovery were explored. Information overload, determining the relevancy of novel knowledge and searching for what “you don’t know you don’t know” were identified as challenges. For the organization, implementation and adoption of innovative ideas – the institutionalization process – was challenging due to the novelty of the idea. Thus, reframing was used as a change management technique to help individuals learn novel concepts. All of these findings provided support for the NKD design theory – both the meta-requirements and proposed design principles.

Individual use of the Athens tool was explored during this case study. Participants found Athens challenging to use due to the novelty of the tool and the task, as well as problems with the tool itself, such as lack of proper-noun filtering. The overall concept of what Athens does was also difficult to grasp since it is very different from existing search tools. Participants described the difficulty of switching from their familiar search behaviours – looking for something specific – to less-familiar behaviours – divergent thinking and making connections between indirectly-related areas.

Technological frames (Orlikowski and Gash, 1994) and adaptive structuration theory (AST) (Desanctis and Poole, 1994) can help shed light on participants’ use of Athens and related challenges. AST describes the structural potential of a technology as the structural features and the “spirit” of these features (Desanctis and Poole, 1994). In the context of the NKD tool –

Athens – the instantiated design principles of the tool represent the structural features; the goals and values of the NKD tool represent the spirit. Outcomes may be unintended or suboptimal if the spirit is incoherent or if appropriation of the features is unfaithful. In this study, it is possible that the spirit of the NKD tool was incoherent to the participants. Although the researcher spent considerable time over the course of the case-study visit discussing the spirit of the NKD tool, there was some degree of confusion amongst participants over what the NKD tool was doing. However, there was evidence that people developed an understanding of the tool's functionality and capabilities as the sessions progressed. In the final wrap-up group session, a discussion was held that helped clarify any misunderstandings regarding the tool's functionality and the need to use Athens in an exploratory manner.

Unfaithful appropriation of the tool is another possible explanation for the observed tool use and outcomes. Other structures, such as the task, and factors, such as individual knowledge and experience, can affect the appropriation of the tool's features (Desanctis and Poole, 1994). In the case of the NKD tool, participants' technological frames regarding Web-search tools were strongly influenced by their experience with Google. These technological frames, strongly associated with focused-search based tools such as Google, and the novelty of the task itself affected how participants appropriated the NKD tool. Some participants were clearly trying to use the Athens tool as a focused-search tool. Thus, learning to use and appropriating the tool features was affected by technological frames and resulted in changing the task from a level 3 to a level 1 task. Some participants used and evaluated the NKD tool as a focused-search tool, which represents an unfaithful appropriation. Since the features of Athens are incongruent with the participants' technological frames, performance with the Athens tool was low. Given the limited amount of experience with the NKD tool, adaptation through the continued interaction with the technology, task and other relevant social structures was not possible in this case study. Longer-term use and adaptation of the NKD tool would probably yield interesting insights regarding novel-knowledge discovery and is an interesting avenue for future research

Despite these challenges, participants were intrigued with the tool and could see its potential for the organization. In addition to the suggestions made for improving the tool, individuals made suggestions for how to use the tool in the organization and the types of sources best suited for novel-knowledge discovery. Less-structured sources with more full text may be better sources for novel-knowledge discovery. The use of Athens as an alerts application was suggested. For example, this idea involved running Athens periodically, with only the new⁸ novel results being sent to individuals as an alert. In addition, participants discussed integrating the system with their domain-specific thesaurus to provide additional guidance regarding important terms. Finally, participants also suggested using the tool internally to discover novel knowledge within their internal knowledge sources.

Insights into the learning processes associated with novel-knowledge discovery were also gained during this study. Participants found search-term definition challenging for novel-knowledge discovery because of the uncertainty involved – “you don’t know what you don’t know”. It was also challenging to define search terms without expertise in the topic area. Information foraging involves multiple iterations of refinement and filtering. Individuals are used to being able to iterate, refine and filter their results during search activities. In most search tools, this iterative adjustment process is manually performed by the individual. However, for novel-knowledge discovery this process is automated by the Athens tool. The lengthy Athens processing time constrains manual refinement of search terms. People are used to focused-search tools where results are provided instantly and direct control over the adjustment and refinement processes possible. Thus, Athens is very different from focused-search tools in regards to information foraging.

Analysis of the tool-use transcripts helped elaborate the construct definitions of intuiting and interpreting processes. Intuiting processes involved pattern recognition and the development of novel insights. Ideas are not yet formed, but are subtle, fuzzy, initial inklings of interesting

⁸ New since the last query processed.

possibilities. While using the tool, individual intuiting processes involved the formation of initial thoughts about points of interest and directions to follow. Individuals sought to clarify, understand, and reduce the uncertainty with this initial thinking. Interpreting involved taking these novel insights and translating them into ideas for the organization to pursue. This translation process, as described by participants during the interview sessions, involved sense making and resulted in updated mental models. Whether the individual views the results using his or her expert intuition – exploitive, based on current mental models – or entrepreneurial intuition – exploratory, based on future possibilities (Crossan et al., 1999) – affected how the individual interpreted and, thus, the usefulness of the Athens results. Most participants interpreted the results using their expert intuition, rather than using the tool in an exploratory manner. Integrating these learning-process insights with the tool issues identified above led to suggestions for tool enhancements to better support intuiting and interpreting processes.

Finally, at the group level, insights were gained into how social networks can contribute to novel-knowledge discovery. Group interpretation processes are important for shaping and developing shared interpretations of ideas. Integrating processes such as pilots and other forms of evaluation are used as part of the “grass roots” method of idea development and implementation.

The results of this case study were compared and triangulated with the results of the lab experiment – Study 1. The qualitative and exploratory analysis conducted in this case study helped provide rich explanations for the findings from Study 1. As a result, updates to the NKD design theory were made and another iteration of the build-evaluate cycle completed. In fact, two iterations of the build-evaluate cycle were completed – the evaluation conducted in the initial visit and the secondary evaluation based on the updated Athens use process. During these evaluations, important insights were gained about filtering proper nouns, which led to modifications of the tool and use processes.

The findings of this study have implications for the knowledge-discovery theoretical framework developed in Chapter 3. Most participants described their search behaviours as level

1 search – focused search. Several different subtypes of focused search were identified. Some level 2 search behaviour was also described – scanning. Despite its importance to the organization, there does not appear to be any intentional level 3 behaviour – novel-knowledge discovery. Thus, one practical implication from these results is that organizations who value novel knowledge should explore how to best incorporate processes and tools to support novel-knowledge discovery in their organization, rather than relying on pure serendipity.

In addition to the insights developed in this study, there are some potential limitations that require further discussion. The first concern is the extent of researcher participation and intervention in this case study. Some may argue that the high level of researcher participation and intervention during both the tool-use sessions and the secondary Athens evaluation represents an action-research design. The original purpose of the tool-use sessions was to passively observe participants as they used the NKD tool. However, the overwhelming novelty of the tool and the task made it necessary for the researcher to more actively assist participants during these sessions. The researcher helped participants navigate through the tool and answered questions about the results. The secondary Athens evaluation was also not included in the original research design, but an emergent opportunity to engage in another build-evaluate cycle. Although the degree of researcher participation and intervention is higher than typically observed in case studies, this research design does not follow the classic five-phase action-research process (Susman and Evered, 1978), nor does it exemplify the other common action-research characteristics, such as an action and change orientation (Peters and Robinson, 1984). The tool evaluation and the build-evaluate cycles included in this case study are characteristics of the design-science research approach adopted in this thesis.

The degree of researcher participation in the case study also introduces a threat to the validity of the results. The researcher's involvement may have caused participants to react to the NKD tool in an overly positive manner. While this validity threat is possible, the results provide evidence that it is unlikely. The results above suggest that participants were willing to be candid

and open with the researcher. Participants provided the researcher with considerable constructive feedback regarding the tool. As depicted in the results, there were more “Athens problems” described by the participants than “positive tool experiences”. Some participants were very candid in their observation that the tool would not be useful in their role, which revolved around focused and comprehensive searching. If an overly positive bias was present, these more critical comments would not have arisen.

Another concern is the objectivity of the researcher in reporting the results of the tool evaluation. While the objectivity of the researcher cannot be completely assured, the inclusion of critical comments about the tool provides evidence that the researcher is not simply reporting positive results. The results discussed above indicate that there were a number of problems and issues with using the Athens tool. Participants found the tool difficult to learn and use. In many instances, the results of the first evaluation were perceived as irrelevant. However, there were some positive reactions to the tool. Most importantly, the participants viewed the tool as quite valuable. As a result, the organization suggested a secondary evaluation and communicated a desire to pursue the development of the tool further. Lastly, the organizational sponsor reviewed this chapter in order to validate the portrayal of the results.

8.7 Summary

In this chapter, the results of a single, revelatory case study were discussed. The purpose of this study was to explore novel-knowledge discovery in practice. Insights were gained into the organizational challenges, processes and tools associated with novel-knowledge discovery. In addition, this case study provided a rich understanding of the multi-level organizational learning processes in the context of novel-knowledge discovery. Two iterations of the build-evaluate cycle were completed in this study, resulting in modifications to the NKD tool and use process.

In Chapter 9, the group processes associated with novel-knowledge discovery are explored further and the results of Study 3 – a field experiment to examine the individual and group processes associated with NKD tool use and outcomes – are discussed.

CHAPTER 9

STUDY 3 – FIELD EXPERIMENT

9.1 Purpose

The purpose of Study 3 – field experiment – was to evaluate different processes of using an NKD tool in the multi-level organizational learning process. Developing an understanding of NKD tool-use processes creates a second design artifact for the NKD design theory. Specifically, Study 3 examines the individual and group processes associated with NKD tool use discussed in Chapter 5 and the outcomes of those processes.

The field experiment addresses propositions 6 and 7 and provides additional validation of propositions 2c, 2f, 3c and 5 (see Chapter 6 for a summary). The main research hypotheses evaluated in this study are as follows:

Radical and Benefit Outcomes

H11: NKD tools combined with individual and collaborative-filtering processes will result in the choice of ideas with higher perceived levels of **“radicalness”** than individual and group processes alone.

H12: NKD tools combined with individual and group processes, such as consensus-building, will result in the choice of ideas with higher perceived levels of **“benefit”** than individual or collaborative-filtering processes alone.

Satisfaction and Confidence Outcomes

H13: NKD tools combined with individual and group processes, such as consensus-building and process support, will result in higher perceived levels of **“satisfaction with the process”** than individual or collaborative-filtering processes alone.

H14: NKD tools combined with individual and group processes, such as consensus-building and process support, will result in higher perceived levels of **“confidence with the decision made”** than individual or collaborative-filtering processes alone.

In the following sections, the research design employed for this field experiment, as well as an overview of the data analysis procedures and results, are discussed.

9.2 Research Design

As described in Chapter 6, this field experiment uses a quasi-experimental design. Specifically, a within-subjects control, with pretest design was employed. One organizational group participated in all three treatment conditions (described below), creating within-subjects control groups. Although random assignment was not used, the groups being contrasted across treatment conditions are equivalent since the same participants completed each treatment condition. Thus, selection bias was not an issue in this experiment.

The design of this experiment incorporated feedback and results from the first two studies. For example, Athens training and instructions were updated to account for the challenges expressed by individuals in both the lab experiment and the case study. In addition, the Athens tool itself was updated to include partially-automated proper-noun filtering, based on feedback from the case study.

In the following sections, the measures, pretest activities and components of the experiment are discussed.

9.2.1 Measures

The constructs measured during the field experiment are described below, including the independent variable, dependent variable and other experimental variables.

9.2.1.1 Independent Variable

The independent variable – treatment condition – measured in this field experiment was the type of process used to generate and choose an idea associated with a level 3 task – novel-knowledge discovery. Three different treatment conditions were used. All conditions included individual use of the NKD tool to search for and generate novel ideas. The difference between conditions and, thus, the manipulation was the process used to evaluate and choose an idea for the organization to pursue. Each treatment condition is described below.

The first process – individual use of the NKD tool and individual evaluation – involved individuals using the NKD tool to search for and generate novel ideas to address the level 3 task. Each individual evaluated his or her own list of ideas and selected the best idea to address the task from this list.

The second process – individual use of the NKD tool plus use of a collaborative-filtering tool for evaluation – involved individuals using the NKD tool to search for and generate novel ideas to address the level 3 task. In this process, the ideas generated by each individual were combined into one list and provided to the group to evaluate individually, without group interaction. The resulting scores were provided to a collaborative-filtering tool in order to “automatically” choose the best idea. The main difference between process 1 and process 2 was the combination of ideas and the use of the collaborative-filtering tool.

Finally, the third process – individual use of the NKD tool plus use of a GSS tool for evaluation – involved individuals using the NKD tool to search for and generate novel ideas to address the level 3 task. Similar to the second process, the ideas generated by each individual were combined into one list and provided to the group to evaluate individually, without group interaction. The scores were combined and the ideas – content and tabulated scores – were discussed among the group to build consensus and clarify uncertainties. After the group discussion, individuals evaluated the ideas individually and anonymously a second time. The results of this evaluation were immediately tabulated and shown to the group, spurring further discussion and a final decision on which idea to pursue. The main differences between processes 2 and 3 were the group discussion resulting in consensus as well as the different information-processing mechanisms. In process 2, the collaborative-filtering algorithm provided the information-processing mechanism, whereas in process 3, a spreadsheet displaying a decision model was the information-processing mechanism used.

9.2.1.2 Dependent Variables

The key dependent variables included: 1) perceived radicalness of the idea, 2) perceived potential benefit of the idea, 3) satisfaction with the process, and 4) confidence in the decision made. In addition, the perceived quality of the idea – how well the idea addresses the task – was measured and referred to as the “overall” score. This measure was used in lieu of ranking¹. As discussed in Chapter 7, new scales were developed to measure the radical and benefit constructs.

The type of knowledge discovered, learning mode and individual learning processes supported were also assessed in order to further validate the findings from Study 1.

9.2.1.3 Other Variables

In addition to the variables mentioned above, a number of other variables were measured during pretest and posttest. To enable a comparison of results, this experiment measured the same pretest and posttest variables measured in the lab experiment. The pretest and posttest instruments are discussed in more detail below.

9.2.2 Experiment Pretest

A number of pretest activities were performed in order to refine and test all aspects of the field experiment. Since there was considerable overlap between the constructs measured in the lab experiment and this experiment, only minor modifications were made to the pretest and posttest instruments. All aspects of the experiment, including the survey instruments, were pretested iteratively using two PhD students, one of whom was an expert in experimental design, and the other had extensive industry experience. In addition, the overall experiment process was reviewed with two members of the researcher’s thesis committee.

Three pretests of the experiment were conducted in the Queen’s Executive Decision Centre (QEDC) lab at Queen’s University. The organization that participated in this experiment

¹ In the initial research design, individuals were to be asked to rank their ideas. However, after pretesting this design, it was decided to use an “overall” rating of the ideas instead of ranking.

had group decision-making facilities similar to those of the QEDC and requested that the experiment be conducted in their own facilities. Thus, another partial “dry run” of the experiment was completed in the organization’s facilities prior to the actual experiment. The experimental process and instruments were updated based on the feedback provided in all pretest activities.

9.2.3 Experiment Execution

In the following sections, details regarding the sample, training, pretest, experimental task and posttest are reviewed.

9.2.3.1 Sample

The organization that participated in this experiment was a Canadian subsidiary of a large Fortune 500 organization in the high-tech industry. Initially, a group consisting of eight individuals was formed and agreed to participate. However, due to unforeseeable conflicts one participant was not able to attend and another individual had to leave part way through the actual experiment. As a result, the final group participating in the experiment consisted of six members. The participants represented a range of roles and departments. A summary of the user profiles is provided below in Table 9.1.

Table 9.1 – Group Profile

User ID	Role	Age Range	Department ²
User 1	Analyst	20-30	Market Intelligence
User 2	Director	40-50	Software Services
User 3	Technical Lead and Creative Director	40-50	User Technology and Design
User 4	Director	50-60	Design Leadership - Strategy and Technology
User 5	Manager	40-50	University Relations
User 6	Senior Director	40-50	Design Research Centre

² Names of departments have been altered slightly to protect anonymity of the organization.

9.2.3.2 Training

A training session was conducted at the organizational site two weeks prior to the experiment. After participants reviewed the letter of information and signed the consent form, the training session began. During the training session, participants were provided with an introduction to novel-knowledge discovery, Athens and the experimental task. Training on how to choose descriptive terms for Athens as well as how to navigate through the Athens results was provided. Based on feedback and results from the case study and the lab experiment, the Athens training included a hands-on component in which individuals were provided with sample Athens results and allowed to interact with those results directly. The researcher answered questions that arose. The training script can be found in Appendix S.

At the end of the training session, participants were provided with a link to and asked to complete the pretest Web questionnaire. An important component of this questionnaire was the provision of Athens descriptive terms. While completing this activity, participants asked the researcher questions about their descriptive terms and refined them accordingly. Some participants provided an extensive list of terms. Similar to the lab experiment, the researcher ran Athens queries for each user based on the terms provided in their pretest questionnaire.

9.2.3.3 Pretest

As noted above, the pretest was administered to participants via Web survey during the training session (see Appendix T for samples of all questionnaires used during the field experiment). The content of the pretest questionnaire was very similar to the one used in the lab experiment. Questions included demographic questions, such as age, gender, tenure, department and role in the organization, familiarity with Web-search tools, familiarity with the subject of the task, Web-specific self-efficacy and general computer self-efficacy. In addition, details regarding the task were provided and associated descriptive terms for the Athens query requested. The pretest questionnaire was estimated to take 10 to 15 minutes to complete. Some participants

spent more time completing their questionnaire because they gave more thought to the generation of their descriptive terms.

During the two weeks between the completion of the training session and pretest, the researcher ran queries on the updated version of Athens. Most participants provided several descriptive terms for Athens. As a result, the researcher ran two queries for every participant. In some cases, the descriptive terms used by participants were the same or similar. In those cases, both participants were given the same Athens results for that query.

The version of Athens used to run these queries was updated to include partially-automated proper-noun filtering. This functionality was added based on the findings from the case study. A description of this filtering mechanism and the process used for running the query is described below.

Each query was first run in Google to ensure sufficient results would be retrieved by the Athens tool to enable the generation of novel clusters. This was a new step employed by the researcher, based on her progressive experience with and understanding of the tool's functionality. If sufficient results were retrieved in the Google search, the researcher initiated the Athens query. During the early processing stage of each Athens query, a dialog box was displayed listing the most prominent proper nouns found during the initial Athens search. The purpose of this dialog box was to allow the user to choose whether any of the prominent proper nouns listed should be included in the proceeding Athens search phases or specifically filtered out. The researcher reviewed each query's proper noun dialog box and used a set of guidelines³ to determine which nouns to keep and which to filter out for the proceeding search phases. These guidelines involved exclusion of the names of people, places or publications from the search. The researcher reviewed the final results of all queries. If necessary⁴, the search terms

³ To ensure a standard process across queries.

⁴ For example, if there was an error or insufficient results.

were modified and the process repeated. The results of the queries were provided to participants at the start of the experimental task.

9.2.3.4 Task

Prior to the start of the experiment, a sheet with the participant's name and descriptive terms was placed at specific computer terminals in the organization's group decision lab. Thus, participants were assigned to specific computer terminals in the lab. On each computer the researcher opened the participant's specific Athens search results – two different queries – and the Web form used to capture ideas during the experimental task.

At the start of the session, participants were provided with a review of the instructions on completing the task and using Athens, both previously discussed during the training session. Supporting handouts for the task and the Athens tool were also distributed. A detailed script was prepared and used by the researcher to ensure consistency of procedure and instructions across all treatments and to ensure the training review was consistent with the original training session (see Appendix U for a copy of the script). The review of the task and Athens tool took approximately 10 to 15 minutes to complete.

After this review, participants were provided with 30 to 40 minutes to complete the task – a novel-knowledge discovery task (level 3 task). The subject of the experimental task was determined by the organizational sponsor in collaboration with the researcher⁵. The task involved identifying novel ways for the organization to generate business value from virtual worlds (see Appendix V for a full description of the task). Participants were asked to generate four to five ideas to address this task using their Athens results to help – either reuse, modify or create new ideas. Participants were asked to document their ideas and the Web link that spurred this idea (i.e., stimulus) in the Web form provided.

⁵ The researcher reviewed the subject with the sponsor to ensure it was appropriate for the experiment.

9.2.3.5 Posttest

After completing the task – search and generation of novel ideas using Athens – the treatment conditions were initiated progressively across three phases of the experiment. The posttest – initiated at the end of each phase – included rating the radicalness and potential benefit of each idea generated during the task, as well as how well the idea addressed the task overall. The progressive nature of the experiment, as well as the relationship between the experiment phases and the processes being compared, are depicted below in Figure 9.1. As shown below, process 1 was represented by phase 1 of the experiment; process 2 was represented by phases 1 and 2; and process 3 was represented by phases 1, 2 and 3.

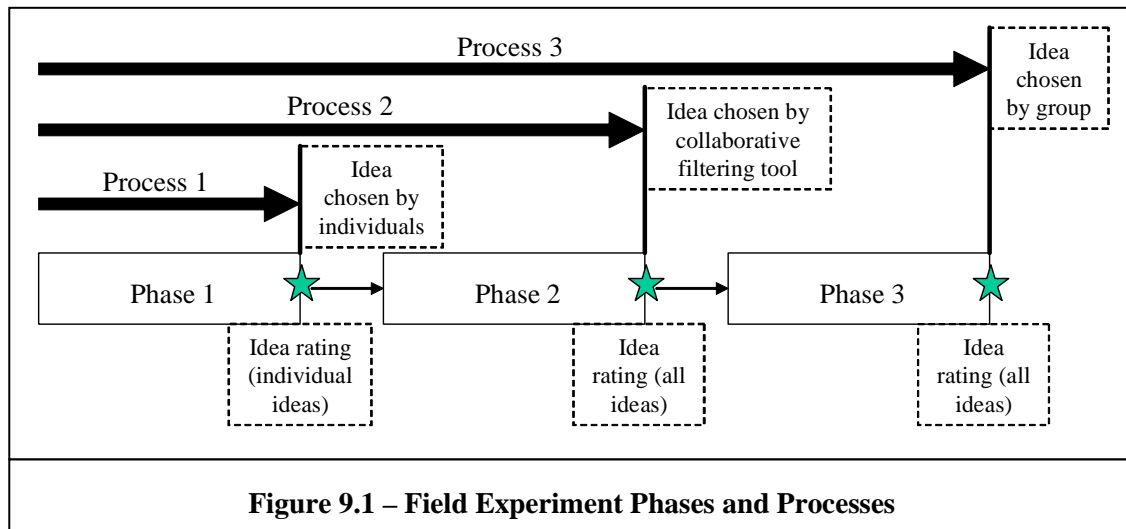


Figure 9.1 – Field Experiment Phases and Processes

Rating was completed using GroupSystems, which was used in the lab experiment. The posttest also included a questionnaire in which participants were asked to indicate their level of satisfaction with the process and confidence in their rating decisions. In addition to these standard posttests, participants were asked to complete a questionnaire at the end of phase 1 in which they were asked to describe the idea they identified as best – both the idea content and rationale – and answer a number of questions about their experience using the tool to complete the experimental task. The questionnaire was almost identical to the one used in the lab experiment and included the knowledge discovery type, learning mode, learning process and PIIT

measurement scale questions. In addition, questions about the task, knowledge of virtual worlds, and experimental control questions were asked. Lastly, participants were asked several qualitative questions including:

- 1) What features of this tool assisted you with today's task?
- 2) What challenges did you encounter using this tool?
- 3) If this tool could be improved, what are the most important features that should be added or fixed?

While using the Athens tool was not the manipulation in this experiment⁶, including these questions about the tool-use experience enabled a comparison of results across studies.

9.3 Analysis

Data analysis involved identifying the top ideas from each process and comparing the radicalness and benefit scores of the top ideas across the processes in order to evaluate the research hypotheses H11 and H12. In process 1, the top idea chosen by each participant was described in the posttest questionnaire completed at the end of phase 1. The top two ideas from process 2 were identified by reviewing the results from the collaborative-filtering tool. Finally, the top two ideas from process 3 were identified by and agreed upon by the participants at the end of phase 3.

Choice shifts between the three phases were analyzed at the individual and group level. In this study, choice shift represented changes in the scores assigned to an idea by an individual and by the group on average.

The level of satisfaction with the decision-making process and the confidence in the decision were compared across the three processes to test hypotheses H13 and H14. Comments made by participants in the posttest surveys and debriefing session were also reviewed and triangulated with the quantitative results. A broader qualitative analysis was also performed,

⁶ The process of evaluating and choosing ideas was the manipulated variable.

including a review of the comments made in the posttest and debriefing session, to help enrich the understanding of the processes and use of the tool.

Finally, the results across the three studies – lab experiment, case study, field experiment – were compared using MANOVA to determine whether there were any significantly different results across studies.

9.4 Results

The results of the field experiment are discussed in detail below, including an overview of the results of each phase of the experiment and comparison of the top ideas across processes. In addition, choice shifts between phases are analyzed and the process satisfaction and decision confidence across processes are compared.

9.4.1 Overview across Phases

In the section below, the results of each phase of the experiment are discussed. The focus of this discussion is on the process followed by the participants during the phase.

9.4.1.1 Phase One

In phase 1 – representing process 1 – participants spent approximately 40 minutes generating ideas based on the results of their Athens queries. During this time, the researcher monitored the progress of the participants in terms of the number of ideas generated⁷. In general, there were three stages noted during this process. In the first stage, participants reviewed the tool results, documenting very few, if any, ideas in the Web form. In the second stage, participants continued to look through their Athens results, and documented the majority of their ideas. In the last phase, participants continued to review their results, but very few ideas were generated on average. Thus, a sequence of slow, rapid, and slow rate of idea generation was noted.

⁷ Using the online survey tool.

Most participants reviewed the results of both of their queries. Some participants commented during the debriefing session that they switched to their second set of results because the first results they reviewed did not provide enough variation or were too novel. One participant commented “I couldn’t make the leap” and understand the connections Athens was making.

At the end of idea generation, participants were asked to rate their ideas on the radicalness, potential benefit and how well the idea addressed the task overall. The question wording for each of these dimensions is outlined in Table 9.2.

Table 9.2 – Instructions for Idea Rating

Dimension	Question
Radicalness	The radicalness of an idea is defined as the extent to which the technical content of the idea is new to the organization, based on what the organization does now. Technical content refers to products, services, processes, materials, technologies. On a scale of 1 to 10, please indicate how radical this idea is to <organization name>, where 1 represents "Not new to <organization name>" and 10 represents "Radically new to <organization name>".
Benefit	The overall benefit of an idea is defined as the potential benefits less potential costs, including tangible and intangible benefits and costs. On a scale of 1 to 10, please indicate how beneficial this idea is to <organization name>, where 1 represents "Not beneficial to <organization name>" and 10 represents "Extremely beneficial to <organization name>".
Overall	In this activity, you will be asked to provide an overall rating for each of your ideas. On a scale of 1 to 10, please indicate how well this idea addresses the search task, where 1 represents "Not well at all" and 10 represents "Extremely well".

This process helped individuals identify which idea they felt was the best idea to propose forward. Participants were asked to complete a Web survey (see Appendix T), in which they were asked to describe the idea they identified as best – both the idea content and rationale – and answer questions regarding satisfaction with the process and confidence in their decision.

The top ideas – chosen by the participants in their individual posttest questionnaires – were ideas 13, 2, 17, 9, 5, and 7. Three of the six participants picked ideas that were balanced in terms of radicalness and benefit scores – either no difference between radicalness and benefit scores, or only a single-point difference. Two participants chose ideas that had higher radical

scores than benefit scores – two-point difference. One participant chose an idea that had a much higher benefit score than radicalness score – five-point difference. None of the participants chose an idea that had a much higher radicalness score than benefit score. Thus, individuals differed in the type of idea they choose as the best idea in terms of benefit versus radicalness.

Phase 1 lasted approximately 90 minutes. Once participants finished completing the posttest questionnaire – signaling the end of phase 1 – a short break was provided while the researcher prepared for the initiation of phase 2.

9.4.1.2 Phase Two

In phase 2 – representing process 2⁸ – participants were provided with an electronic document that contained a description, as well as the associated Web link, for every idea generated during phase 1. There were 21 ideas in total. All ideas were anonymous and sorted randomly. Participants spent 20 to 30 minutes reviewing the ideas and accessing Web links to review the content that spurred each idea⁹. There was no interaction between participants during this phase. Thus, each participant developed his or her own interpretation of the idea. At the end of this review, participants were asked to rate all ideas using the same criteria as phase 1 – radicalness, benefit and overall scores. The rating scores were fed into a collaborative-filtering tool. This processing was performed as a post-experiment analysis. Thus, the results of the collaborative filtering were not available to participants for review.

Once the idea evaluation was completed, participants were asked to complete a brief questionnaire to indicate their level of satisfaction with the process and confidence in their rating decisions. The completion of the questionnaire signaled the end of phase 2. Participants were provided with a short break, during which time the researcher tabulated the results in preparation for phase 3.

⁸ Process 2 includes phases 1 and 2.

⁹ Accessing the Web link was described as optional.

9.4.1.3 Phase Three

The goal of the final phase – representing process 3¹⁰ – was to reach a consensus regarding the best two ideas. At the start of this phase, the researcher presented a decision model that included the tabulated results of the idea evaluation from phase 2 (i.e., democratic votes). The decision model included the mean and standard deviation of the radical, benefit and overall scores for each idea. In addition, the difference between the mean radical and benefit scores for each idea were presented. To aid the review of this data, the results were colour-coded to highlight top scores, large standard deviations and large differences between radical and benefit scores. Table 9.3 provides a sample of the data that was displayed to participants at the start of phase 3.

Table 9.3 – Preliminary Decision Model – Phase 3¹¹

Idea #	Overall	Std Dev	Radical	Std Dev	Benefit	Std Dev	Benefit-Radical Differences
13	9.00	0.89	6.67	1.97	8.67	1.21	2.00
9	8.00	2.10	7.83	1.72	7.17	2.79	-0.67
4	7.83	1.60	8.00	2.45	8.00	1.26	0.00
12	7.83	0.41	6.83	3.31	7.50	1.38	0.67
6	7.67	2.25	7.17	1.83	7.67	1.51	0.50
14	7.67	1.21	7.67	1.75	6.50	2.51	-1.17
7	7.50	1.52	7.67	1.75	7.67	1.51	0.00
18	7.50	2.74	7.00	1.79	7.33	1.51	0.33
21	7.17	2.48	7.83	1.72	7.50	2.07	-0.33
2	7.17	0.75	7.33	3.08	7.17	1.17	-0.17
20	6.83	3.25	6.33	3.20	6.83	2.40	0.50
1	6.67	1.37	7.50	3.08	5.67	2.42	-1.83
3	6.67	1.86	7.33	2.80	5.67	1.51	-1.67
11	6.67	2.58	6.50	2.81	6.17	2.86	-0.33
10	6.33	1.86	8.50	1.52	5.67	1.63	-2.83
19	6.33	2.16	7.83	0.75	5.33	1.86	-2.50
16	5.67	2.50	8.33	1.37	6.50	2.26	-1.83
8	5.50	1.52	7.67	1.21	6.00	2.19	-1.67
15	5.00	2.19	6.33	1.37	5.33	1.37	-1.00
5	4.50	2.07	6.50	2.07	6.83	1.17	0.33
17	4.33	2.50	9.00	1.10	3.50	2.07	-5.50

¹⁰ Process 3 includes phases 1, 2 and 3.

¹¹ Idea summary was also displayed, but removed from this display to protect confidentiality.

Facilitated by the researcher, participants reviewed the scores and discussed the content of the ideas in an effort to come to a shared understanding of the ideas. In some cases, participants discussed the rationale behind why they felt an idea was beneficial or radical or why it was problematic. During this time, participants began clustering certain ideas that represented a common theme. Over the course of the discussion, participants discussed 12 out of the 21 ideas (ideas 1, 2, 3, 4, 9, 11, 12, 13, 17, 18, 19, 21).

Participants also discussed the measures used to rate the ideas. It became clear that the participants had different perspectives on the radical and benefit measures. One participant indicated they viewed radicalness negatively. This participant commented, “High radical is not necessarily a good thing”. Another participant responded,

“But if we want to look for some measure of innovation. So higher on radical but still fits the business model. I figured the combination of radical and high business value were good. So a new opportunity different from what anybody is doing, but would actually generate cash.”

Participants recognized they viewed and, therefore, measured radicalness differently. They also noted that the definition of radical is relative to the individual’s level of knowledge – what is radical to one person may not be radical to another based on their knowledge. Before the next round of evaluations, participants agreed upon defining radicalness as “innovativeness to the organization”.

After this discussion, participants were asked to anonymously rate all of the ideas a second time. Upon completion of this rating, the researcher tabulated the results and displayed the updated decision model to the participants (see Table 9.4).

Table 9.4 – Final Decision Model – Phase 3

Idea #	Overall	Std Dev	Radical	Std Dev	Benefit	Std Dev	Benefit-Radical Differences
4	8.67	1.63	6.67	3.88	6.17	3.19	-0.50
6	8.67	1.21	6.00	1.67	5.83	3.54	-0.17
9	8.50	1.87	7.67	2.94	6.83	3.19	-0.83
13	7.83	3.43	4.33	2.58	9.00	1.26	4.67
18	7.67	1.86	5.33	3.33	6.50	3.15	1.17
7	7.17	2.14	7.50	3.08	7.17	2.48	-0.33
2	7.17	1.83	6.17	2.32	5.17	2.23	-1.00
12	7.17	1.33	4.50	3.02	5.83	2.64	1.33
21	6.83	2.14	5.67	2.42	6.67	2.58	1.00
20	6.33	3.44	4.83	3.31	4.83	3.49	0.00
16	5.83	3.13	7.00	2.10	5.33	2.88	-1.67
19	5.50	2.59	6.33	2.25	4.17	3.13	-2.17
11	5.00	3.52	4.50	3.45	4.33	3.01	-0.17
14	5.00	3.29	5.67	2.73	5.50	2.66	-0.17
15	3.83	1.60	2.67	1.86	3.17	2.04	0.50
1	3.83	3.13	5.33	3.01	3.33	2.25	-2.00
5	3.67	1.75	4.00	2.10	4.17	1.72	0.17
8	3.67	1.75	5.17	3.06	3.17	2.14	-2.00
10	3.67	2.66	7.67	2.58	2.83	2.14	-4.83
3	3.33	2.73	7.33	2.42	2.67	1.97	-4.67
17	2.83	2.04	8.33	1.86	3.00	1.79	-5.33

After the second evaluation, a facilitated discussion was held in which the group chose the best two ideas. The new results were compared to the previous evaluation results. During this discussion, participants focused on the top overall scores in order to reach a consensus regarding the top two ideas. For one of the top ideas, participants clustered ideas together. Initially, they clustered ideas 9, 7 and 6 together to form one of their top ideas. While finalizing their decision, one participant suggested adding idea 18 to this cluster. The group agreed. This cluster of ideas was viewed as large in scale and costly to implement. However, the group felt this idea could be quite beneficial to the company. The idea was a natural progression for the organization and wasn't viewed as particularly radical. One participant commented, "I don't see it as that innovative. Is that a problem?" Another participant responded, "No. These are things we do well."

The other idea chosen to be amongst the top two ideas was idea 4. This idea was viewed as smaller in scale than the first idea, but much more innovative. One participant commented, “I think it’s (the idea) critical to the future of virtual worlds”.

Once these ideas were chosen, participants were asked to complete a brief questionnaire to indicate their level of satisfaction with the process, confidence in their rating decisions and which decision making process – phase 1, 2 or 3 – they preferred and why. This signaled the completion of phase 3.

After phase 3 was completed, lunch and a debriefing session were held to discuss the experiment. The experiment was approximately three and a half hours in duration, not including breaks.

9.4.2 Chosen Ideas

In order to evaluate hypotheses H11 and H12, the radicalness and benefit scores of the ideas chosen as best in each of the three processes were compared. The results of this comparison are described below.

9.4.2.1 Process 1 – Individuals

As discussed above, six ideas were chosen in process 1 – one for each participant. In order to compare the results across processes, the collective ratings for these ideas from the voting in phase 2 are reported below (see Table 9.5)¹².

Table 9.5 – Process 1 Top Ideas – Results

Idea	Radical	SD*	Benefit	SD*
13	6.67	1.97	8.67	1.21
2	7.33	3.08	7.17	1.17
17	9.00	1.10	3.50	2.07
9	7.83	1.72	7.17	2.79
5	6.50	2.07	6.83	1.17
7	7.67	1.75	7.67	1.51
Average	7.50		6.83	

*Standard deviation

¹² Using scores from each phase would confound the comparison of ratings of top ideas with choice shift.

9.4.2.2 Process 2 – Individuals and Collaborative Filtering

For process 2, the overall, radical and benefit ratings were entered into the collaborative-filtering tool. The collaborative-filtering tool uses matrix decomposition – specifically singular value decomposition (SVD) – to compute a global assessment of idea ratings, weighted by each participant’s ability to evaluate. The “quality” of an individual assessment is based on the degree to which an individual’s opinions are in agreement with the opinions of other individuals with better judgment skills (Skillicorn, 2001). Collaborative filtering resolves this circular relationship and provides a collective recommendation, biased towards strong assessors. As discussed in Chapter 5, one can think of a strong assessor as a “discerning” evaluator. Using movie ratings as an illustrative example, the collaborative-filtering tool distinguishes between two types of raters: 1) raters who show an ability to discern different “grades” of movies, similar to other “high quality” raters, and 2) raters who like or dislike most movies. “High quality” raters are assigned more weight so that the global recommendation is skewed towards the movies, or ideas in this case, that these raters evaluate highly. In theory, the results from a collaborative-filtering tool should differ from a democratic vote.

The phase 2 votes served as input to the collaborative-filtering tool in the form of a matrix – idea ratings (row) by individual rater (column). Matrices for each of the three dimensions – overall, radical and benefit – were created and fed into the collaborative-filtering tool separately. However, the top ideas for process 2 were identified using the “overall” results since participants were told to think of the overall score as how well the idea addressed the task, which included “novelty” and “value” as key elements. The raw output consisted of three matrices referred to as U, S, and V in singular value decomposition. This raw output was transformed into a series of two- and three-dimensional graphs to facilitate interpretation and analysis of the results.

To determine the ranking of ideas, the output of the collaborative-filtering tool was interpreted by, first, analyzing the most prominent singular values. Singular values represent underlying factors or dimensions in the data (Skillicorn, 2007). The number of prominent singular values was determined by reviewing the diagonal values in the S matrix. These values can be interpreted in a similar fashion to eigenvalues in factor analysis. A review of the values in the S matrix suggested two prominent singular values or dimensions. These dimensions were interpreted by reviewing the graphical output of the tool, which included global ratings of ideas and ratings of individuals. Reviewing the graphical output of the global idea ratings, the two prominent dimensions were interpreted as: 1) overall quality of the idea, and 2) breadth of the idea in terms of application.

Once meaningful dimensions were identified, the collaborative-filtering tool was used to draw a line that best represented these two factors simultaneously and project the ideas onto this line. The output of this processing was a list of the ideas and values indicating positioning on the line. The position of each idea on the line represented a “rank” order of ideas. Some interpretation was required to determine the most appropriate sort direction (i.e. descending vs. ascending).

The idea rankings produced by the collaborative-filtering tool for the overall score are shown below (Table 9.6). The collaborative-filtering tool produced global-idea rankings that differed from the democratic-vote rankings (for democratic results see Table 9.3 above).

Table 9.6 – Collaborative Filtering Idea Ranking

Idea Ranking	Position
9	-4.20
18	-3.07
13	-2.73
16	-2.65
4	-1.71
12	-1.15
10	-0.77
7	-0.57
21	-0.53

Idea Ranking	Position
6	-0.49
11	-0.40
14	-0.13
2	0.37
19	0.46
20	1.33
15	1.53
17	1.63
1	2.42
8	2.70
5	3.52
3	4.46

The top ideas identified by the collaborative-filtering tool and their radical and benefit scores are presented in Table 9.7. Since process 1 and process 3 each included 6 and 5 ideas respectively, the top 5 ideas were analyzed for process 2.

Table 9.7 – Process 2 Top Ideas – Results

Ideas	Radical	SD*	Benefit	SD*
9	7.83	1.72	7.17	2.79
18	7.00	1.79	7.33	1.51
13	6.67	1.97	8.67	1.21
16	8.33	1.37	6.50	2.26
4	8.00	2.45	8.00	1.26
Average	7.57		7.53	

*Standard deviation

9.4.2.3 Process 3 – Individuals and Group Support System

For the third process, the participants came to a consensus regarding the top two ideas at the end of phase 3. As noted above, one of these ideas was a cluster of four related ideas. For comparability, the collective ratings for these ideas from phase 2 voting are reported below.

Table 9.8 – Process 3 Top Ideas – Results

Ideas	Radical	SD*	Benefit	SD*
9**	7.83	1.72	7.17	2.79
7**	7.67	1.75	7.67	1.51
6**	7.17	1.83	7.67	1.51
18**	7.00	1.79	7.33	1.51
4	8.00	2.45	8.00	1.26
Average	7.53		7.57	

*Standard deviation

**Four ideas forming a cluster

A summary of the average radical and benefit scores for the top ideas chosen in each process are presented in Table 9.9.

Table 9.9 – Summary of Top Ideas across Processes

Process	Radical	Benefit
Process 1 (6 ideas)	7.50	6.83
Process 2 (5 ideas)	7.57	7.53
Process 3 (5 ideas)	7.53	7.57

These results suggest that the average radicalness of the ideas chosen in each process does not vary across processes. The average radicalness score is very similar across all three processes. Thus, H11 can be rejected. The average benefit score for process 3 is larger than the score for process 1. However, process 3 and process 2 have very similar average benefit scores. Thus, H12 is partially supported. It is important to recognize that process 2 is automated and requires much less work than process 3. Given the similarity in scores between processes 2 and 3, efficiency is an important factor to consider when choosing evaluation processes.

Given the small sample size, a Kruskal-Wallis test was conducted to assess whether there were statistically significant differences between the three processes in terms of radicalness or benefit scores. The results are presented below in Table 9.10. There were no statistically significant differences between processes in terms of radicalness or benefit.

Table 9.10 – Evaluation of Differences between Processes – Radicalness and Benefit

	Process 1	Process 2	Process 3	Test Statistics
Radical				
Mean	7.50	7.57	7.53	
Standard Deviation	0.91	0.70	0.43	
Kruskal-Wallis (mean rank)	7.67	9.30	8.70	Chi-Square = 0.338
				p = 0.845
Benefit				
Mean	6.83	7.40	7.57	
Standard Deviation	1.75	0.79	0.32	
Kruskal-Wallis (mean rank)	6.92	9.00	9.90	Chi-Square = 1.181
				p = 0.554

An interesting observation about the ratings and chosen ideas across processes is that in process 3 participants appeared to shy away from purely beneficial ideas, for example idea 13. Idea 13 was originally the top idea listed in the preliminary decision model (see Table 9.3). However, after the group discussion, this idea was rated lower overall and had a much higher standard deviation (i.e. less agreement). Thus, in the context of this task, individuals rated low ideas that were highly beneficial, but not radical.

9.4.3 Choice Shift

As discussed in Chapter 5, individuals in a group tend to shift their preferences and ratings of alternatives after a group discussion. One reason for these shifts is partially-shared persuasive arguments, which are brought up during the group discussion (Vinokur, 1969; Vinokur and Burnstein, 1974). In order to examine these shifts in preferences and ratings, choice shifts were analyzed at the individual level between phases 1 and 2, and between phases 2 and 3. Choice shift was also analyzed at a group level – average scores for each idea – between phases 2 and 3¹³. The results of the individual level and group level analyses of choice shift are discussed in more detail below.

¹³ Phase 1 only included votes by individuals on their own ideas, and thus wasn't included in the group level analysis of choice shift.

9.4.3.1 Individual-Level Choice Shifts

Analysis of choice shifts at the individual level indicated that participants shifted their individual scores between phases 1 and 2, tending more towards positive shifts than negative shifts (see Tables 9.11 through 9.13 below). Both the average shift and the average absolute shifts were calculated for the overall, radical and benefit scores.

Table 9.11 – Overall Ratings: Individual Choice Shift between Phases 1 and 2

Idea	User1	User2	User3	User4	User5	User6
1	2					
2		1				
3			3			
4				3		
5					1	
6						0
7						0
8					3	
9				2		
10			-1			
11		0				
12	3					
13	1					
14						0
15					2	
16				2		
17			1			
18	-1					
19						1
20	2					
21						1
Average Absolute Shift	1.80	0.50	1.67	2.33	2.00	0.40
Average Shift	1.4	0.5	1	2.3	2	0.4
	positive shift					
	negative shift					

Table 9.12 – Radical Ratings: Individual Choice Shift between Phases 1 and 2

Idea	User1	User2	User3	User4	User5	User6
1	0					
2		0				
3			-2			
4				1		
5					-2	
6						3
7						3
8					0	
9				1		
10			5			
11		1				
12	0					
13	0					
14						3
15					-1	
16				2		
17			1			
18	1					
19						3
20	-1					
21						1
Average Absolute Shift	0.40	0.50	2.67	1.33	1.00	2.60
Average Shift	0	0.5	1.3	1.33	-1	2.6
	positive shift					
	negative shift					

Table 9.13 – Benefit Ratings: Individual Choice Shift between Phases 1 and 2

Idea	User1	User2	User3	User4	User5	User6
1	1					
2		-1				
3			0			
4				3		
5					-2	
6						0
7						0
8					2	
9				3		
10			-2			
11		0				
12	0					
13	1					
14						-1
15					1	
16				2		

Idea	User1	User2	User3	User4	User5	User6
17			2			
18	-2					
19						-3
20	3					
21						0
Average Absolute Shift	1.40	0.50	1.33	2.67	1.67	0.80
Average Shift	0.6	-0.5	0	2.7	0.3	-0.8
	positive shift					
	negative shift					

In contrast to the individual shifts between phases 1 and 2, the shifts between phases 2 and 3 are mostly negative. Further, the shifts between phases 2 and 3 appear to be much larger in magnitude than the shifts between phases 1 and 2 (see Tables 9.14 through 9.16 below).

Table 9.14 – Overall Ratings: Individual Choice Shift between Phases 2 and 3

Idea	User1	User2	User3	User4	User5	User6
1	0	0	1	-4	-7	-7
2	3	1	-1	-2	-2	1
3	-1	-1	-2	-4	-8	-4
4	1	1	1	2	1	-1
5	1	2	1	-5	-3	-1
6	0	4	-1	2	0	1
7	0	2	-2	-3	0	1
8	0	1	-1	-5	-3	-3
9	1	1	1	0	1	-1
10	-1	-1	0	-7	-5	-2
11	0	0	3	-6	-6	-1
12	0	1	0	-2	-3	0
13	0	0	1	1	-9	0
14	-1	-1	-1	-6	-6	-1
15	-1	1	3	-6	-3	-1
16	1	2	2	1	-4	-1
17	2	2	-3	-6	-2	-2
18	1	3	-1	-3	1	0
19	1	2	0	-1	-8	1
20	0	-3	5	-6	0	1
21	0	1	1	-1	-5	2
Average Absolute Shift	0.71	1.43	1.48	3.48	3.67	1.52
Average Shift	0.3	0.86	0.3	-2.9	-3.4	-0.9
	positive shift					
	negative shift					

Table 9.15 – Radical Ratings: Individual Choice Shift between Phases 2 and 3

Idea	User1	User2	User3	User4	User5	User6
1	0	-4	2	-1	-5	-5
2	0	-2	2	0	-5	-2
3	0	-3	2	-2	3	0
4	0	-5	1	0	-5	1
5	-3	-6	1	-5	-2	0
6	1	-4	-3	-3	1	1
7	0	-7	-2	1	5	2
8	-3	-6	-3	-5	3	-1
9	1	-6	-2	0	2	4
10	0	-5	-2	-3	4	1
11	0	-4	0	-6	-2	0
12	-2	-6	-3	-3	0	0
13	0	-5	-2	-3	-4	0
14	-1	-2	-2	-2	-4	-1
15	-1	-6	-2	-5	-6	-2
16	1	-4	-3	0	-1	-1
17	0	-2	1	-3	2	-2
18	0	-6	-4	-5	3	2
19	1	-2	-3	-2	-3	0
20	1	-5	-2	-2	0	-1
21	0	-5	-4	0	-4	0
Average Absolute Shift	0.71	4.52	2.19	2.43	3.05	1.24
Average Shift	-0.2	-4.5	-1.3	-2.3	-0.9	-0.2
	positive shift					
	negative shift					

Table 9.16 – Benefit Ratings: Individual Choice Shift between Phases 2 and 3

Idea	User1	User2	User3	User4	User5	User6
1	-3	-5	2	-3	-2	-3
2	-1	-2	-1	-2	-4	-2
3	0	-4	-2	-5	-5	-2
4	0	-6	-1	2	-3	-3
5	-1	-2	-2	-6	-2	-3
6	0	-2	-3	-5	1	-2
7	2	-3	-3	-2	1	2
8	1	-3	-2	-5	-4	-4
9	0	-2	-1	0	0	1
10	-1	-2	0	-7	-4	-3
11	-1	-2	-3	-3	-1	-1
12	-1	-1	-1	1	-7	-1
13	0	0	-1	1	0	2
14	0	-1	1	-2	-4	0
15	1	-2	-2	-4	-3	-3

Idea	User1	User2	User3	User4	User5	User6
16	-1	-3	-1	0	0	-2
17	0	-2	-2	-1	3	-1
18	2	-4	-1	-3	-4	5
19	0	-2	2	0	-4	-3
20	0	-3	3	-3	-7	-2
21	0	-3	0	0	1	-3
Average Absolute Shift	0.71	2.57	1.62	2.62	2.86	2.29
Average Shift	-0.1	-2.6	-0.9	-2.2	-2.3	-1.3
	positive shift					
	negative shift					

In order to assess whether the shifts between phases 1 and 2, and phases 2 and 3 differed significantly, a Kruskal-Wallis test was performed on the individual average shifts for each dimension – overall, radical and benefit. Significant differences were found between the shifts in phases 1 and 2, and phases 2 and 3 for all three dimensions (α 0.01 and 0.05 levels). The shifts from phase 2 to phase 3 were negative on average, whereas the shifts from phase 1 to phase 2 were positive on average (see Table 9.17 below).

Table 9.17 – Differences between Choice Shifts – Individual Average Shift

	First shift	Second shift	Test Statistics
Overall			
Mean	1.27	-0.96	
Standard Deviation	0.78	1.80	
Kruskal-Wallis (mean rank)	9.17	3.83	Chi-Square = 6.587
			p = 0.010**
Radical			
Mean	0.79	-1.57	
Standard Deviation	1.24	1.64	
Kruskal-Wallis (mean rank)	9.00	4.00	Chi-Square = 5.789
			p = 0.016*
Benefit			
Mean	0.38	-1.57	
Standard Deviation	1.24	0.97	
Kruskal-Wallis (mean rank)	9.17	3.83	Chi-Square = 6.564
			p = 0.010**

** significant at the 0.01 level

* significant at the 0.05 level

9.4.3.2 Group-Level Choice Shifts

At the group level, choice shifts between phases 2 and 3 were analyzed. The literature that discusses group choice shifts suggests that partially-shared persuasive arguments brought up during group discussion may result in choice shifts based on the content of the argument. The implication is that after the group discussion the opinions and, thus, ratings of ideas across individuals should be more similar than before the group discussion. During the group discussion in phase 3, participants discussed the rating criteria in an effort to develop a shared understanding of the criteria. Thus, any differences in ratings due to variation in the meaning of rating criteria should have been reduced or eliminated. On several occasions participants shared their rationale for liking or disliking a particular idea – possibly a partially-shared persuasive argument. An example of such a dialogue follows,

“I figured the combination of radical and high business value were good. So new opportunity different from what anybody is doing, but would actually generate cash. I liked idea 13 – had business value and was most elegant. There were lots of great ideas about what to do in virtual worlds, but in terms of generating cash...”

“I found that one the hardest to rate.”

“It’s just a good idea about what might happen in virtual worlds. But we don’t have any stake in it so there may be no benefit to <organization name>.”

Participants also commented on the fact that the group discussion helped develop shared understandings of ideas and rating criteria. This shared understanding should have led to a greater degree of agreement on ratings.

Contrary to this prediction, the standard deviation in scores for each idea increased from phase 2 to phase 3. Table 9.18 depicts the results of the group-level choice shift analysis. On average, participants decreased their ratings of ideas – more so for the radical and benefit dimensions. The decrease in ratings may be the result of participants taking a more critical perspective of ideas as the group attempted to move closer to choosing an idea for the

organization to pursue. The arguments made during the discussion may have helped participants become more critical. However, the level of agreement in ratings actually decreased signaling lack of consensus and shared understanding. Thus, it's possible that some participants were more critical of ideas than others.

Table 9.18 – Group-Level Choice Shift

Idea	Overall	Standard Deviation	Radical	Standard Deviation	Benefit	Standard Deviation	AVERAGE SHIFT
1	-2.83	1.76	-2.17	-0.07	-2.33	-0.17	-2.44
2	0.00	1.08	-1.17	-0.76	-2.00	1.06	-1.06
3	-3.33	0.87	0.00	-0.38	-3.00	0.46	-2.11
4*	0.83	0.03	-1.33	1.43	-1.83	1.92	-0.78
5	-0.83	-0.32	-2.50	0.02	-2.67	0.55	-2.00
6*	1.00	-1.04	-1.17	-0.16	-1.83	2.04	-0.67
7*	-0.33	0.62	-0.17	1.33	-0.50	0.98	-0.33
8	-1.83	0.23	-2.50	1.85	-2.83	-0.05	-2.39
9*	0.50	-0.23	-0.17	1.22	-0.33	0.40	0.00
10	-2.67	0.80	-0.83	1.07	-2.83	0.50	-2.11
11	-1.67	0.94	-2.00	0.64	-1.83	0.15	-1.83
12	-0.67	0.92	-2.33	-0.29	-1.67	1.26	-1.56
13	-1.17	2.54	-2.33	0.62	0.33	0.05	-1.06
14	-2.67	2.08	-2.00	0.98	-1.00	0.15	-1.89
15	-1.17	-0.59	-3.67	0.50	-2.17	0.67	-2.33
16	0.17	0.62	-1.33	0.73	-1.17	0.62	-0.78
17	-1.50	-0.46	-0.67	0.77	-0.50	-0.28	-0.89
18*	0.17	-0.88	-1.67	1.54	-0.83	1.64	-0.78
19	-0.83	0.43	-1.50	1.50	-1.17	1.26	-1.17
20	-0.50	0.19	-1.50	0.11	-2.00	1.09	-1.33
21	-0.33	-0.35	-2.17	0.70	-0.83	0.51	-1.11
AVERAGE SHIFT	-0.94	0.44	-1.58	0.63	-1.57	0.71	

*Top ideas chosen by the group

Legend

	Score	Standard Deviation
	Positive shift (group voted more favorably on average).	Difference in ratings increased.
	Small negative shift (group voted less favorably on average, by less than a point).	Difference in ratings decreased, by one point or less.
	Large negative shift (group voted less favorably on average, by more than a point).	Difference in ratings decreased, by more than one point.

In order to assess whether the ideas that were chosen as the top ideas had significantly different group-level choice shifts from the rest of the ideas, a Kruskal-Wallis test was performed. The results of this analysis are depicted below in Table 9.19. There were significant differences between the chosen and non-chosen ideas across all dimensions (α 0.01 through 0.10 levels). The chosen ideas had, on average, an increase in overall scores whereas the non-chosen ideas had a decrease in overall scores (α 0.01 level). In addition, the overall-score standard deviation decreased for the chosen ideas, signifying consensus, compared to the increased standard deviation for non-chosen ideas (α 0.05 level). For the radical and benefit scores, the decrease in ratings for the chosen ideas was significantly less than for the non-chosen ideas (α 0.10 level). However, the standard deviation for the radical and benefit scores increased for the chosen ideas more than for the non-chosen ideas (α 0.10 and 0.05 level). Thus, from an overall score perspective, the ideas that were ultimately chosen at the end of phase 3 showed more consensus and favorable scores than the non-chosen ideas. This same result was not carried over to the radical and benefit scores.

Table 9.19 – Differences between Group Choice Shifts

	Non-chosen Ideas	Chosen Ideas	Test Statistics
Overall Shift			
Mean	-1.36	0.43	
Standard Deviation	1.06	0.53	
Kruskal-Wallis (mean rank)	8.69	18.40	Chi-Square = 9.364
			p = 0.002**
Overall Standard Deviation Shift			
Mean	0.67	-0.30	
Standard Deviation	0.91	0.68	
Kruskal-Wallis (mean rank)	12.56	6.00	Chi-Square = 4.261
			p = 0.039*
Radical Shift			
Mean	-1.79	-0.90	
Standard Deviation	0.88	0.69	
Kruskal-Wallis (mean rank)	9.56	15.60	Chi-Square = 3.626
			p = 0.057 [†]

	Non-chosen Ideas	Chosen Ideas	Test Statistics
Radical Standard Deviation Shift			
Mean	0.50	1.07	
Standard Deviation	0.70	0.70	
Kruskal-Wallis (mean rank)	9.75	15.00	Chi-Square = 2.727
			p = 0.099 [†]
Benefit Shift			
Mean	-1.73	-1.07	
Standard Deviation	0.94	0.72	
Kruskal-Wallis (mean rank)	9.75	15.00	Chi-Square = 2.743
			p = 0.098 [†]
Benefit Standard Deviation Shift			
Mean	0.49	1.40	
Standard Deviation	0.50	0.69	
Kruskal-Wallis (mean rank)	9.38	16.20	Chi-Square = 4.609
			p = 0.032*

** significant at the 0.01 level

* significant at the 0.05 level

† significant at the 0.10 level

9.4.4 Process Satisfaction and Decision Confidence

In order to determine whether there were differences between processes in the level of satisfaction with the process (H13) and confidence in the decision (H14), a Kruskal-Wallis test was performed, followed by Mann-Whitney tests for post-hoc analysis of differences. The results are presented below in Table 9.20.

Table 9.20 – Evaluation of Differences between Processes – Satisfaction and Confidence

	Process 1	Process 2	Process 3	Test Statistics
Satisfaction				
Mean	5.00	6.00	6.17	
Standard Deviation	0.63	0.89	0.75	
Kruskal-Wallis (mean rank)	5.42	11.00	12.08	Chi-Square = 6.016
				p = 0.049*
Mann-Whitney (mean rank)	4.67	8.33		p = 0.058 [†]
Mann-Whitney (mean rank)	4.25		8.75	p = 0.022*
Mann-Whitney (mean rank)		6.17	6.83	p = 0.733
Confidence				
Mean	5.00	6.00	6.33	
Standard Deviation	0.89	1.10	0.82	

	Process 1	Process 2	Process 3	Test Statistics
Kruskal-Wallis (mean rank)	5.67	10.67	12.17	Chi-Square = 5.340
				p = 0.069 [†]
Mann-Whitney (mean rank)	4.83	8.17		p = 0.093 [†]
Mann-Whitney (mean rank)	4.33		8.67	p = 0.031*
Mann-Whitney (mean rank)		6.00	7.00	p = 0.604

* significant at the 0.05 level

[†] significant at the 0.10 level

Significant differences were found between processes for satisfaction with the process (α 0.05 level). As predicted, participants were more satisfied with process 3 than process 1 (α 0.05 level). Although satisfaction with process 3 was not significantly higher than process 2, the responses to the question “Of the three phases in this experiment, please indicate which phase you preferred most in terms of the decision-making process” provided evidence that participants were more satisfied with process 3 than processes 1 and 2. Of the six participants, five indicated they preferred process 3 and one indicated they preferred process 1. Comments made by participants in the posttest survey and during the debriefing session further supported the preference and higher satisfaction with the third process. In the posttest survey two participants commented on process 3:

“Taking the group input and honing it into a synthesized result was most useful as well as seeing the bounds on the ratings using the standard deviations.”

“I thought the collaboration aspect was the critical one, where we discussed pros and cons of each idea, and then came to consensus.”

In the debriefing session two participants commented,

“I thought the last phase was most critical for me. I liked the face-to-face collaboration that the tool didn’t provide at all. It was nice to go over all our thoughts and flush out some of the ideas and rescore. It just seems more valuable that way.”

“There’s a difference between doing something purely individually vs. online (in a group). Where we discovered – for example – the term radically was being used radically differently. You either continue that confusion all the way through or you get it clarified or even the description of items.”

Thus, the quantitative and qualitative data provided support for H13 – higher levels of satisfaction for process 3.

Significant differences were also found between processes for confidence with the decision made (α 0.10 level). Participants were more confident in decisions made in process 3 than process 1 (α 0.05 level). Although process 3 had a higher mean score and rank than process 2, the difference was not significant. Thus, H14 was partially supported.

9.4.5 Participant Comments – Process and Tool

The process-related comments in the posttest questionnaire and in the debriefing session highlighted some potential future process improvements regarding the use of an NKD tool to generate and choose ideas to pursue. One participant suggested getting group agreement on the meaning of the evaluation criteria prior to rating the ideas. Although definitions were provided for the criteria, allowing participants to “personalize” and align the criteria to the organization’s context was suggested. In addition, this participant suggested including a phase where participants cluster the ideas and evaluate those clusters, rather than individual ideas. This participant commented,

“I would have liked to have more collaboration / consensus building along the way. It would have helped if we agreed on definitions of 'radical', for example, and if we were able to cluster our ideas or differentiate them better. This could be done keeping anonymity if we did it using the GroupSystems capabilities.”

Further, several participants suggested including idea ranking in addition to rating ideas. However, participants felt that rating ideas on multiple criteria such as radicalness and benefit was useful. Participants commented,

“It was good to have different dimensions. Helps you refine your reasons and think through why the idea is good or not. Ranking might have been better.”

“I agree multiple criteria is important.”

“Rate first and then rank the shortlist.”

Thus, using both idea rating followed by idea ranking of the top five or ten ideas was suggested as a process improvement.

Comments about the use of Athens to generate ideas were also made by several participants. Individuals used the tool in different ways to generate novel ideas. One participant found that the ongoing process of reviewing and progressing through the Athens results helped her to generate ideas, not just the review of a single Web link. Another participant depended less on the specific results, but used the results as a jumping off point to think more creatively. One participant commented,

“Does seem to be very individual. Some people liked to look at the result set, which by itself spurs off ideas. Then there are people like <name removed> who wanders off on his own, who is naturally curious. All depends on how we use it. I used it as a tool to assist versus a tool that is going to give you results.”

Another participant found that even the descriptive nouns helped stimulate ideas. Thus, scanning the high-level results, without going into detail, can help spur ideas. This participant felt that a tool like Athens could be useful in her job.

Comments from several participants suggested that an NKD tool like Athens could help groups with brainstorming tasks. One participant commented on the speed with which the group came up with several innovative ideas for consideration through the individual and group processes. In this participant’s experience, this type of process typically lasts for days, rather than half a day in this case¹⁴. Another participant indicated that they found the limited amount of time to generate ideas using the tool helpful. She noted that it usually takes her a long time to generate ideas and find information on the Web. Thus, she found that the tool and process really helped speed this process up and facilitate information foraging, idea generation and stimulate thinking. It is plausible that the indirectly-related results helped stimulate creativity. This finding is supported by the group-creativity literature (Hender et al., 2002; Nagasundaram and Bostrom, 1995).

In addition to the tool being used in different ways by different people, several participants commented on the usefulness and interpretation of the results being dependent on the

¹⁴ Both processes used the organization’s decision lab and GroupSystems software.

individual's existing knowledge-base. Thus, the concept of absorptive capacity and how it impacts the use of Athens was discussed. One participant also noted that Athens doesn't provide new information. Rather, it provides a new view of existing information by reframing how you look at that information and connect it to the task or subject area. This comment provides evidence that the participants – at least this particular participant – understood the functionality and goal of Athens (i.e., the “spirit” of the features (Desanctis and Poole, 1994)).

The type of question and knowledge area under consideration were noted as important inputs to the process when using an NKD tool. Several participants questioned whether the group had selected the right question. They felt that the emerging area of virtual worlds may not have been an appropriate subject for a novel-knowledge discovery task – the topic was novel itself. One participant commented,

“The topic is important. I worried about our topic in that there wasn't much about it period. And as a result we got really novel connections. There just wasn't enough content for Athens to work on.”

Thus, some emerging topics – novel themselves – may not be good candidates for novel-knowledge discovery on the Web. However, additional research is needed to further examine the relationship of topic to the effectiveness of an NKD tool.

9.4.6 Comparison to Other Studies

Each of the three studies in this dissertation – lab experiment, case study and field experiment – collected the same posttest data, providing a basis for comparison between studies. Based on the results of a MANOVA test, there were no significant differences found in the quantitative results across the three studies. Interestingly, the group means for the novel, interpretation and learning-mode constructs were higher in the third study than the other studies, but not significant. This could indicate a trend of improvement in the tool's ability to support novel-knowledge discovery and help individuals adjust their mental models.

The most interesting and interpretable comparison points between the three studies were the qualitative comments about the Athens features, challenges and suggestions for improvement. These comments are discussed in more detail below.

9.4.6.1 Features

Similar to the lab experiment, participants in the field experiment mentioned that the clustering functionality of Athens, as well as the ability to drill down into more detail, were particularly helpful. Similar to the case study, participants mentioned that the tag clouds on the novel cluster pages were useful. These features helped participants develop an understanding of the results and choose which clusters they should investigate in further detail. Participants in the field experiment commented that they liked having access to the results of multiple queries. Interestingly, one participant commented that the descriptive nouns from both the intermediate clusters and the novel clusters were helpful in generating novel ideas. Thus, the nouns themselves helped stimulate novel-idea generation in this case.

9.4.6.2 Challenges

Participants in the field experiment experienced many of the same challenges as those found in the lab experiment and the case study. For example, participants sometimes found the results too novel, too similar, or irrelevant. Participants also commented on the lack of Web link metadata. Interestingly, there were fewer comments made about proper-noun filtering issues. Thus, the improvements made to Athens to address the influence of proper nouns helped to resolve this particular problem, which was detected in the earlier studies.

9.4.6.3 Suggested Improvements

Many of the same suggestions made in previous studies regarding improvements to Athens were also made by the participants in this field experiment. For example, suggestions were made regarding improving the interface layout and structure, as well as changing the display

to focus the user's attention on the most interesting results. One participant suggested enabling interactivity with the clusters – for example, bookmarking, grouping clusters and initiating new Athens queries based on interesting clusters selected by the user.

Another interesting similarity across the three studies was found in the manipulation check, which was used to determine whether individuals understood the task they were being asked to perform. Four of the six participants understood the task. However, two participants were looking for specific things in their search task, despite the training instructions asking individuals to be open, exploratory and not to look for specific things. Thus, these individuals appeared to be performing a level 1 task. Similar to the results of the other studies, these individuals found it difficult to break free from their technological frames regarding Web-search tools and searching for specific things.

9.5 Discussion

The process of using an NKD tool to generate and choose a novel idea to pursue was the key area being examined in this experiment. Three different processes and their outcomes were compared in this study, including individual and group processes. The results showed very little difference in the average radicalness of the top ideas chosen in each process. However, ranking the results by the average-radicalness score places process 2 above the other processes – providing limited support for H11. On the benefit dimension, process 3 had the highest average-benefit score for the top ideas. However, process 2 and 3 had very similar average-benefit scores. Thus, H12 was partially supported. As expected, participants were more satisfied with process 3 and more confident in the decision made in this process than the other processes. Thus, there was support for H13 and partial support for H14.

In order to further understand the evaluation of ideas across processes, the choice shifts between phases were compared at the individual and group levels. At the individual level, the results showed that the choice shift between phase 1 and phase 2, in which there was no group

interaction, had a positive shift – more favorable ratings. However, once group interaction was introduced – between phases 2 and 3 – the choice shift was negative, resulting in less favorable ratings on average. Thus, the group discussion influenced how individuals perceived the ideas and resulted in participants taking a more critical perspective.

At the group level, some individuals were more critical than others as indicated by the increase in standard deviations (i.e., less agreement) between phases 2 and 3. Interestingly, the ideas that were ultimately chosen after the final vote actually had a positive shift – more favorable ratings – and lower standard deviations in scores (i.e., more agreement). Thus, the group had come to a shared understanding of these ideas. Based on this shared understanding, the group was able to come to a consensus on choosing these ideas. While this finding is interesting, it brings into question why the group only came to a shared understanding of a few ideas – 5 out of 21. Part of this may be the result of time restrictions. However, Boland and Tenkasi (1995) suggest that a groupware voting system, where finding a consensus is promoted, can hamper “the team members from first strengthening and representing their own perspectives and then engaging in a dialogue of perspective taking with each other” (Boland and Tenkasi, 1995 p. 360). Thus, the emphasis on consensus and time restrictions may have affected the results of this study.

Comments made by participants during this study reveal that group discussion – perspective taking – is critical for developing a shared understanding of novel ideas, part of the group interpreting process in the 4I organizational learning process model (Crossan et al., 1999). Evaluating those ideas and making a decision regarding which idea to pursue – group integrating processes – is also dependent on group discussion. The limited number of ideas in which a shared understanding was gained in process 3 may be the result of the novelty of ideas and associated cognitive complexity. In the context of novel-knowledge discovery, groups may be limited in the number of ideas they can consider due to the novelty of ideas. Further, groups may require additional discussion time in order to fully discuss and reach a shared understanding of all ideas.

During the course of the experiment, several process improvements were suggested for the overall process of generating and choosing novel ideas to pursue. First, developing a shared understanding of the measures being used – also referred to as personalizing the measures – was deemed an important step. In addition, including idea ranking of the top five ideas in the evaluation process was suggested.

The radicalness and benefit scores of the top ideas did not differ significantly between processes 2 and 3. Given process 2 requires less effort than process 3, organizations may want to consider using a collaborative-filtering tool in place of lengthy group decision meetings. A blend or combination of the two processes may be more appropriate given the higher satisfaction and confidence scores for process 3. Combining these two processes is discussed below.

A comparison of the chosen ideas across the three processes shows that the set of top ideas chosen in each phase does indeed differ to some degree. The collaborative filter used in process 2 produced results that differed from both the democratic vote in phase 2 and the consensus reached in process 3. Thus, for novel idea evaluation, using a collaborative-filtering tool as a decision model and comparing it to a democratic vote may help highlight important ideas that may not have come out in the discussion. For example, ideas 13, 16 and 12 – within the top 6 ideas for process 2 – were not explicitly chosen in process 3. Thus, using a collaborative filter as a decision model could help identify ideas that may be worth discussing further during the group interaction phase of process 3.

This field experiment enabled the development of an extended and more in-depth understanding of how an NKD tool can support novel-knowledge discovery and learning in organizations. In this study, it was evident that individuals used the NKD tool in different ways and, as a result, differed in their opinions of how useful the tool was for the task. As discovered in the other studies, some individuals found a level 3 task – novel-knowledge discovery – difficult. These individuals found making the “leap” between the subject of the task and the novel results presented by the NKD tool difficult. Part of this difficulty may be due to the prototypical

state of the Athens tool. However, the difficulties encountered may also be due to the influence of technological frames. For example, some individuals had expectations that the tool would behave like other search tools in terms of query response time and types of results provided. Further, some individuals were clearly looking for something specific – a level 1 task – rather than being more exploratory in their search. Thus, overcoming technological frames is a major challenge for novel-knowledge discovery.

Despite these challenges, there were positive outcomes and evidence that Athens supported novel-knowledge discovery and learning during the completion of the level 3 task. The output of the tool helped individuals generate novel ideas to address the experimental task. For one participant, simply reviewing the descriptive nouns helped her make the “leap” and generate novel ideas. The tool also appeared to provide some degree of support for individual learning processes, such as intuiting and interpreting. The descriptive nouns, tags, clusters and content of the Web pages within the novel clusters helped individuals detect patterns – intuiting – and make sense of the results in order to generate novel ideas – interpreting. For one participant, the ongoing review and exploration of the results contributed to both the process of intuiting and interpreting. The tool also supported mental model building. One participant commented on how much they learned about the subject of the task – virtual worlds – through this experience. Another participant commented on the tool’s ability to reframe one’s perspective of a subject. These results help support the NKD design theory and propositions 2c, 2f, 3c and 5.

The usefulness and potential of an NKD tool was validated through the comments made by the participants. One participant in the experiment commented that a tool like Athens would be useful in her job. In addition, after the completion of the experiment the organizational sponsor of this study sent the researcher a summary of the feedback received from individuals regarding the chosen ideas and the Athens tool. These individuals thought that the chosen ideas were interesting and inquired as to whether anyone was pursuing them in the organization. Further, one of the participants in the study indicated they would like to install the Athens tool on their system

so they could experiment and use it further within her team. This is a strong signal that despite Athens prototypical state and array of current issues, there is potential value in the tool for organizations interested in novel-knowledge discovery and innovation.

Although these results are promising there are some limitations, which impact the generalizability of the results. One of the most significant limitations of this experiment was sample size. Due to the large time commitment required on the part of the participating organization, only one organizational group was included in this experiment. Thus, the results and generalizability may be limited by the small sample size. However, the collection of both quantitative and qualitative data provided rich insights, which can be further explored and validated in future research.

The structured nature of the phases in this experiment, as well as the time restrictions, may have limited the findings of this study. Allowing participants to evaluate clusters rather than individual ideas and providing additional time for discussion may have created a more realistic decision environment and potentially different results. However, the structure was necessary for experimental control and comparability across phases. As a result of busy work schedules, participants had limited time to devote to this experiment. Thus, increasing the duration of this experiment may have resulted in fewer participants and negatively impacted the findings.

Comparing the results across all three studies may be problematic due to the changes made to the NKD tool and use process between these studies. However, comparing new and improved versions of the NKD tool is a valuable part of the design-science process. The design-science process and build-evaluate cycles across the three studies are discussed in more detail in Chapter 10.

Plausible threats to the validity of the results include the treatment diffusion, testing effects, and changes in instrumentation across phases. The progressive nature of the experiment may have resulted in treatment diffusion. However, the progressive design was used to allow for

within-subjects control groups and more comparable results between treatments. The result was a more natural process for the group to follow. Further, the treatment conditions were represented by the cumulative phases, for example process 3 included phases 1, 2, and 3. Therefore, the progressive design was both effective and efficient for this particular study.

Testing effects is another possible threat to validity. Participants were asked to evaluate the ideas a number of times, which may have affected how they evaluated the ideas over time. However, multiple evaluations of the ideas were part of the treatment by design.

During phase 3 the participants discovered they each viewed the measure “radicalness” differently. Before the final vote, participants agreed on a common definition. Thus, changes to the instrumentation between phases 1, 2 and 3 could have affected the results. In order to account for changes in instrumentation, the analysis of choice shift and the comparison of the top ideas across processes were completed separately. The choice shift analysis included the effects from the changes in instrumentation. Despite the changes in instrumentation, the choice shift analysis produced interesting results regarding differences in shifts between “chosen” and “non-chosen” ideas. Further, the comparison of chosen ideas used phase 2 votes in order to avoid confounding the results with choice shift and instrumentation changes.

Future research should examine group and organizational processes associated with novel-knowledge discovery, including the process improvements suggested in this research.

9.6 Summary

Using a field experiment, this study examined how individuals within a group might use an NKD tool – Athens – to accomplish an organizational task requiring novel knowledge. Three different processes of generating and choosing novel ideas to pursue were compared. An understanding of the NKD tool-use processes represents the second NKD design theory artifact. This study supplements the results from the case study and lab experiment. As the third study in this dissertation, this experiment included enhancements to the Athens tool and tool training

based on feedback and learning from the previous studies. Thus, this study represents another iteration of the build-evaluate cycle, which is critical to the design-science process. A comparison of the quantitative and qualitative results across the three studies shows a trend of improvement, suggesting the enhancements made have had a positive impact.

CHAPTER 10

DISCUSSION AND CONCLUSION

A summary of the key findings and insights across all three studies are discussed below. This includes a discussion of how the results contribute to the design science build-evaluate cycle and how they address the two research questions identified earlier. The strengths and contributions of this research, as well as the limitations are also provided. Lastly, future research opportunities are presented.

10.1 Discussion

The research questions identified in Chapter 3 were as follows:

1. What design properties should a novel-knowledge discovery tool have in order to be effective?
2. How do tools oriented towards different levels of knowledge discovery affect the organizational learning process?

A discussion of how these questions were addressed in Studies 1 through 3 is presented below.

10.1.1 Research Questions

In order to address the first research question – “What design properties should a novel-knowledge discovery tool have in order to be effective?” – an NKD design theory was developed in Chapter 4 and assessed across all three studies. In Study 1 – lab experiment – an instantiation of the design theory (Athens) was evaluated and compared to two other tools. The lab experiment helped assess whether the NKD design principles identified in Chapter 4 would be more efficient and effective at supporting novel-knowledge discovery than tools with other design properties. These design principles were developed to address the requirements for novel-knowledge discovery, including principles to address information overload, recognize the

relevance of novel knowledge and learn novel concepts. The design principles were based on learning theories – both cognitive and organizational learning theories.

In Study 2 – case study – participants used the NKD tool Athens and commented on the strengths and importance of the tool, as well as the design characteristics that would make the tool more effective for novel-knowledge discovery. The individual interviews provided rich data regarding the effectiveness of the tool as well as the current challenges, thus further elaborating on the results from Study 1.

In Study 3 – field experiment – a group used the NKD tool Athens to accomplish an organizational task requiring novel knowledge. Two different components of the NKD design theory were evaluated in the third study: 1) the design principles of the tool, and 2) the process of using the tool. Specifically, three different processes were compared, including individual, collaborative-filtering and group processes. The outcomes of these processes were compared.

In all three studies, data were collected regarding the most useful tool features, challenges and suggested improvements. Many similar themes were found across the three studies, strengthening the reliability and validity of the findings. The findings of each study contributed to either process changes or tool changes in the subsequent study. These build-evaluate iterations of design science are discussed in further detail below.

The second research question – “How do tools oriented towards different levels of knowledge discovery affect the organizational learning process?” – was also assessed in all three studies. In the theoretical framework developed in Chapter 3, three different types of tools were proposed. In this framework, it was proposed that each tool most effectively supports the discovery of particular types of knowledge – deep, broad and novel – as well as different learning modes and processes. In Study 1, three different types of tools were used to complete a level 3 task (novel-knowledge discovery task) and compared using an experimental design. In addition to measuring the perceived type of knowledge discovered, the degree of learning support – changes to mental models and learning process – were also measured. In Study 1, hypotheses

associated with three different theories were compared: 1) the theoretical framework and NKD design theory, 2) the competing theory of technological frames, and 3) an integrated theory combining these two theories.

In Studies 2 and 3, the type of knowledge discovered and learning support provided by an NKD tool (level 3 tool) were measured both quantitatively and qualitatively. These results were used to triangulate to and support the results from Study 1. While these studies primarily evaluated individual learning processes, the impact an NKD tool has on group learning processes was assessed in the third study.

Across these three studies, all levels of the 4I organizational learning process (Crossan et al., 1999) were examined. In Study 1, the focus was on individual learning processes. All three levels of “human” learning processes – individual, group and organizational – were explored in Study 2. The focus of Study 3 was on group learning processes. The fourth learning process proposed in this research – machine level – was also assessed across all three studies.

The key findings related to these research questions and three studies are described in the following section.

10.1.2 Key Findings

There were several key findings and insights found in this research. These include the influence of technological frames, the importance and usefulness of an NKD tool to organizations, and insights regarding effective processes of using an NKD tool. These key findings are discussed below.

As discovered across all three studies, technological frames associated with Web-search tools significantly influenced the use of the NKD tool. There was strong evidence that Google’s search interface, proposed to be the dominant design for Web-search tools, had a significant effect on the content of an individual’s technological frames. These technological frames influenced not only how individuals perceived the results of the NKD tool, but the task itself. For

example, some individuals attempted to perform a level 1 task instead of the requested level 3 task. As a result, Google tended to outperform the other tools, providing support for the competing theory – technological frames. However, interestingly the third theory used to evaluate the results of Study 1 provided evidence that despite the strong influence of technological frames, the NKD tool (Athens) does support certain aspects of novel-knowledge discovery and learning more effectively than other tools. Specifically, Athens was more effective at discovering surprisingly new knowledge and reframing mental models. Thus, another important insight was the discovery of multiple dimensions of novel knowledge.

The field studies – case study and field experiment – provided insights regarding the importance of novel knowledge and the usefulness of an NKD tool such as Athens. In Study 2, participants indicated that novel knowledge was vital to innovation and, thus, their organization's survival. However, the discovery of novel knowledge was purely serendipitous in this organization. There was no intentional novel-knowledge discovery and, thus, no tools specifically designed to support the discovery of novel knowledge. In both studies, participants indicated that Athens provided stimulus that allowed them to reframe their thinking in order to learn novel concepts and assess the relevance of those concepts for the organization. This effect was likely due to the indirectly connected results provided by Athens, which are contextually appropriate and linked to the focal topic through an intermediate and directly related topic. Thus, the indirectly connected results provided by NKD tools such as Athens differ from those of tools designed to promote creative thinking through the provision of randomly unrelated stimulus. The efficiency of novel-knowledge discovery was also noted as an important benefit of Athens in both field studies. Although Athens was still in a prototypical state with a number of implementation and interface issues, both organizations expressed interest in pursuing the use of Athens in their respective organizations. This is very positive evidence in support of the importance and usefulness of an NKD tool to organizations.

This research also provided insights into how group learning processes – interpreting and integrating – are influenced by novel-knowledge discovery. In Study 2, participants discussed the group processes associated with novel-knowledge discovery in their organization. In Study 3, participants actually performed the group processes of interpreting and integrating as part of a novel-knowledge discovery experimental task. Three different processes were compared in a quasi-experimental design. Although the outcomes of each process did not differ significantly in terms of the radicalness and benefit of the chosen ideas, the actual ideas suggested at the end of each process did differ.

Participants felt more satisfied and confident in the decisions made with the process that involved group discussion. The choice-shift trends clearly identified the favoured ideas before the group reached a consensus¹. Thus, the initial group discussion in this process helped the group develop a shared understanding of the ideas. While this process – referred to as process 3 – is clearly important, the collaborative-filtering process – referred to as process 2 – produced results that differed from both the democratic vote and the consensus reached by the group. Given the cognitive complexity of reaching a shared understanding of novel ideas, it may be important to integrate a collaborative-filtering tool with group-discussion processes. For example, using a collaborative-filtering tool as a decision model and comparing it to a democratic vote may help highlight important ideas that may not have come out in the discussion, but are worth further examination. Further, the use of a collaborative-filtering tool may help make the decision-making process more efficient by automating part of the analysis.

10.1.3 Design-Science Cycle

As described in Chapter 4, there are three core components of design science: IS research, environment and knowledge base (Hevner et al., 2004). A discussion of how each of these components was updated during this research is presented below.

¹ Choice shift was analyzed after the experiment was finished.

10.1.3.1 IS Research

A major part of the IS research component is the build-evaluate cycle – an important aspect of the design-science process (Hevner et al., 2004). Several iterations of the build-evaluate cycle were performed across the three studies (see Figure 10.1). After the NKD design theory was “built”, an early prototype of the NKD tool (Athens 1.0) was evaluated against this theory and gaps identified.

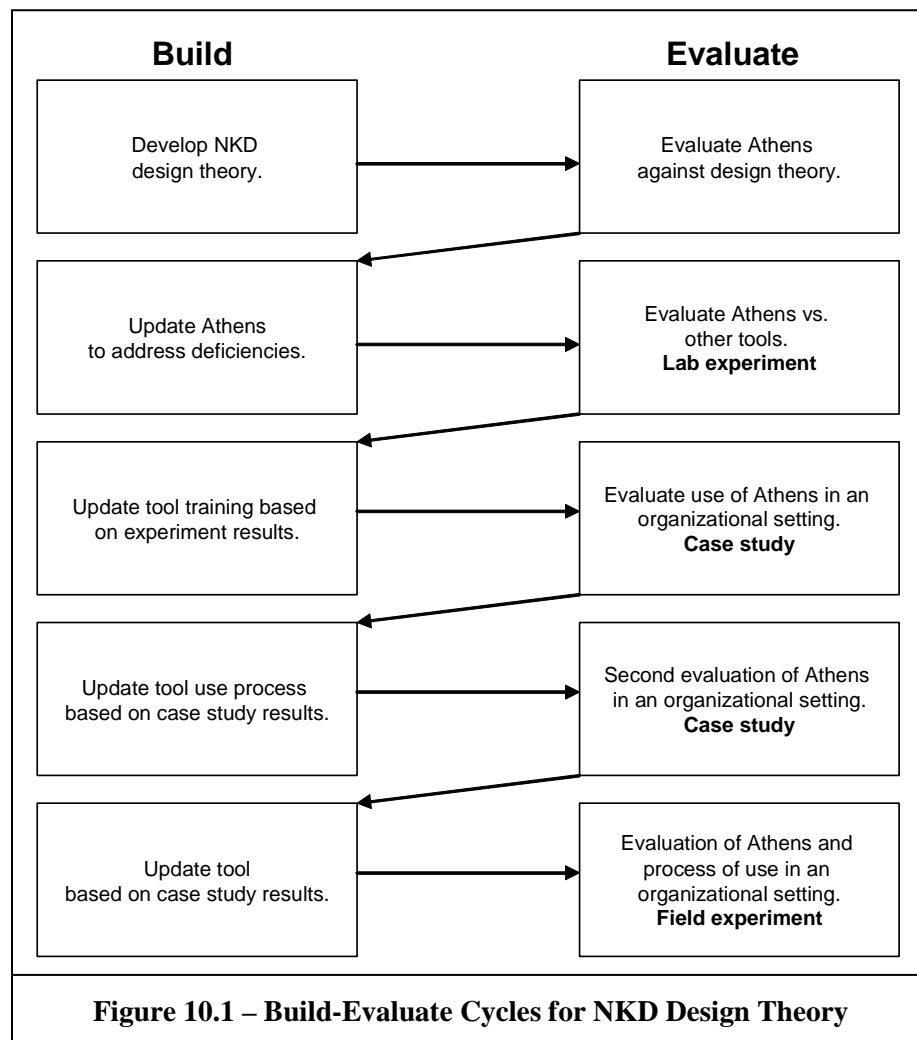
In the second build-evaluate cycle, changes were made to the Athens tool to address these gaps (Athens 2.0) and evaluated in Study 1 – lab experiment – by comparing the effectiveness of the NKD tool to other tools. One of the insights generated from this lab experiment was the strong influence of technological frames and the tendency to use the Athens tool to look for something specific and, thus, perform a level 1 task instead of a level 3 task. Insights were also gained about the challenges and useful features of the tool.

In the third build-evaluate cycle, the training associated with the Athens tool was updated to reflect the insights from Study 1 regarding technological frames. In Study 2 – case study – the differences between Athens and other tools were highlighted and participants instructed to use the tool in an exploratory fashion and not to look for specific things. In addition, other functions of the tool not utilized in Study 1 were examined. For example, the ability of the tool to search a particular Web site or domain was examined. The tool was evaluated by the participants of the case study during one-on-one tool-use observation sessions. Participants still struggled with performing a level 3 task due to their technological frames. However, some progress was made. Further insights into the challenges associated with using the tool were gained. One major insight was the influence of proper nouns on the search.

In the fourth build-evaluate cycle, the process of using the Athens tool was updated to address this proper-noun issue. Prior to the end of Study 2, the researcher reran the Athens queries using this new process – manually filtering the proper nouns using the stopwords list –

and performed a second round of evaluations on the results. Significant improvements in the results were noted by the participants.

In the fifth build-evaluate cycle, the Athens tool was updated to partially automate the filtering of proper nouns. Athens was evaluated in Study 3 – field experiment. A number of insights regarding the tool and the process of using the tool were generated during this evaluation. Thus, both artifacts of the NKD design theory were evaluated in Study 3 – instantiation and tool-use process – and the insights gained can be used for future build-evaluate cycles.



The outcomes of the build-evaluate cycles include an updated design theory, and contributions to both the environment and knowledge base, discussed below.

10.1.3.2 Environment

The environment defines the problem space and includes elements such as people, organization and technology. The field studies conducted during this research – Studies 2 and 3 – provided an enriched understanding of the environment associated with novel-knowledge discovery. For example, both the importance and challenges associated with novel-knowledge discovery in organizations were validated and a deeper understanding developed. Findings from these studies revealed a lack of existing technologies to perform novel-knowledge discovery in organizations. In addition, insights into the different types of people, in terms of roles and tasks, for whom an NKD tool may be useful were gained. Institutionalizing the use of an NKD tool in an organization may require new roles and tasks in order to best leverage the tool and minimize the influence of technological frames. The two major contributions to the environment are the NKD tool Athens, deemed by the field study participants to be a useful and potentially important tool, and insights regarding the use process and outcomes.

10.1.3.3 Knowledge Base

The knowledge base includes theoretical and methodological knowledge that contributes to the development of the design theory. This research has made a number of contributions to the knowledge base related to novel-knowledge discovery. These contributions include the knowledge-discovery theoretical framework, updated organizational learning process, and an understanding of the group processes associated with novel-knowledge discovery, including the use of a collaborative-filtering tool. In addition, instruments for measuring several aspects of the knowledge-discovery theoretical framework were developed and validated. The NKD design theory was also updated based on the findings from Studies 1 through 3.

10.1.3.4 Updated Design Theory

Across the three studies, evidence was found for the adherence of the Athens tool to the NKD design principles. While Athens now adheres to all design principles proposed in Chapter

4, the degree of adherence varies. Certain principles are minimally adhered to and, thus, Athens requires further development and refinement in those areas. For example, the provision of metaknowledge and measurement of interest require additional development in the tool. Participants in all studies recommended using more graphics and visuals to help users understand and make sense of the results. Although Athens is not a perfect instantiation of the NKD design theory, a number of improvements in the tool and its adherence to the principles were made throughout this research.

Based on the results of all three studies, the NKD design theory was updated to include an eighth design principle. This eighth design principle is focused on addressing the challenge associated with technological frames. An NKD tool should be designed in a way that takes into consideration technological frames. For example, including design features that are similar to those of the tool influencing the technological frames – Google in this case. However, the NKD tool also needs to include features that differentiate it from those tools so that individuals use the NKD tool for level 3 tasks and not level 1 tasks. This new design principle, based on the concepts described in Hargadon and Douglas (2001) regarding innovation adoption, is one that could be applied to any novel tool that differs significantly from the dominant design – a meta-meta-design principle.

10.2 Contributions

This research provides a number of contributions to design theory, behavioural theory, practice, and addresses a number of calls for additional research. The contributions in each of these areas are discussed in detail below.

10.2.1 Design Theory

In this research, an NKD design theory was developed to better understand the path to discovering novel knowledge, a significant challenge for organizations. This NKD design theory

can be described as a hermeneutic inquiry system, which supports “human inquiry as a process of subjective, interpretive, meaning making” (Boland et al., 1994 p. 459). This design theory can guide future researchers in the study of novel-knowledge discovery, learning novel concepts, and modifying mental models and assumptions.

Two artifacts were created as part of this design: an instantiation, Athens 2.0, and method of use for the NKD design theory. By using a design-theory approach this research focuses directly on the IT artifact, a critical part of MIS research that requires more attention (Orlikowski and Iacono, 2001).

In addition to the NKD design theory, a new decision model for group decision-making was presented and evaluated in this research. A collaborative-filtering tool, in conjunction with the NKD tool, was evaluated as one of the processes compared in Study 3. The collaborative-filtering tool, which uses singular value decomposition, produced different results from the typical decision-making techniques such as consensus building and democratic vote. Shown to be useful in the NKD context here, collaborative filtering may also be useful in other decision-making contexts and should be examined in greater depth in future research.

10.2.2 Behavioural Theory

In addition to the contributions made to design theory, several contributions to behavioural theory were also made. A knowledge-discovery theoretical framework was developed to identify the sets of tool characteristics proposed to support the discovery of different types of knowledge, as well as support different modes of learning and learning processes. As a result, this paper extends past research that has linked tools to other learning contexts (e.g. Leidner and Jarvenpaa, 1995; Piccoli et al., 2001; Vandenbosch and Higgins, 1996; Vandenbosch and Huff, 1997).

A competing-theories approach was used to evaluate the theoretical framework. Using this approach, a richer and more complete understanding of knowledge discovery was developed

(Robey and Boudreau, 1999; Webster and Watson, 2002). By recognizing the “opposing force” of technological frames and developing an integrated theory, the ability of NKD tools to efficiently support the discovery of surprisingly novel knowledge and reframing were identified. These capabilities would not have surfaced if a competing-theories approach was not used.

This research extends the 4I organizational learning process model (Crossan et al., 1999) specific to the context of searching and learning on the Web. A fifth process is proposed – information foraging and search-term development. In addition, a fourth level is proposed – the machine-level – which describes the level of tool support for individual learning processes such as information foraging, intuiting and interpreting.

New measurement scales were developed and validated in this research for knowledge-discovery types, individual learning processes, and mental model tuning constructs. The exploratory factor analysis conducted in Study 1 produced meaningful factors with satisfactory scale reliabilities. In addition to these knowledge-discovery and learning scales, instruments were developed to measure the potential impact of new ideas in terms of potential benefit and radicalness. Instrument refinement and confirmatory work should be completed in future research.

10.2.3 Practice

There are several contributions to practice made by this research. The NKD design theory can guide practitioners in the development of NKD tools. In addition, organizations can benefit from the use of NKD tools, which can help organizations discover additional strategic opportunities, innovate, or enhance the organization’s learning capability, potentially providing strategic advantage (Zack, 2002; Zack, 2005). Thus, NKD tools support managed learning.

Understanding the influence that technological frames have on the use of NKD tools, as well as other novel tools, can help developers and managers choose the most appropriate users of

NKD tools. Further, training and integration of the tool into routines should be designed to consider the influence of technological frames.

The results from the evaluation of different processes of using an NKD tool to generate and choose an idea in a group setting will help managers understand the outcomes and implications of these processes. Involving multiple individuals in the use of the NKD tool to generate ideas provides diverse input into the decision-making process. Using multiple decision models, such as a collaborative-filtering tool and a democratic-voting model, should help groups develop a shared understanding of a greater number of novel ideas and improve decision making outcomes for the organization.

Lastly, understanding which types of tools might be useful in different learning contexts may assist practitioners in developing their knowledge management strategy and systems, and enable them to “manage” their learning about their external environment (Zack, 2005). For example, firms that are interested in recognizing strategic opportunities arising in the environment may be interested in acquiring or developing level 3 tools. Further, firms could use different types of tools to help diversify the knowledge held across the firm and, thus, increase their absorptive capacity (Cohen and Levinthal, 1990).

10.2.4 Calls for Additional Research

This research addresses a number of calls for additional research in the areas of organizational learning, knowledge management, design science and group support systems. The contributions to each of these areas are discussed below.

By utilizing organizational learning theories, this research addresses calls for more MIS studies in this area (Huber, 1991; Orlikowski and Barley, 2001). Specifically, this research addresses calls for additional research into how different tool characteristics support different aspects of learning (Piccoli et al., 2001) including finer distinctions of knowledge categories (Dutta and Crossan, 2005), how IT can support organizations in the examination and modification

of assumptions, norms and mental models (Boland et al., 1994; Orlikowski and Gash, 1991; Vandenbosch and Higgins, 1996), as well as the exploration of the uncertain relevance of novel knowledge (Schulz, 2001). This research also addresses the need for additional conceptual work on search, information interpretation and cumulative work in organizational learning (Huber, 1991).

By integrating social-science theory with design theory, this research addresses the call for additional design-theory research in MIS to complement the predominant social-science theory research (Hevner et al., 2004; Markus et al., 2002; Venable, 2006; Walls et al., 1992).

This research also contributes to a number of areas within the GSS literature. The need for research that examines more realistic and natural usage of GSS (Fjermestad and Hiltz, 1999) was addressed by the field experiment (Study 3). This field experiment also addressed the call for “more emphasis on the use of larger groups of nonstudent subjects, using more complex tasks over a longer period of time than has been typical” (Fjermestad and Hiltz, 1999 p. 27). For instance, in Study 3 an organizational group performed a complex task involving novel-knowledge discovery, idea generation and decision making over a four-hour time period. Thus, Study 3 contributes to this particular call for research in almost all areas. Further, Study 3 examined the creativity of the group-decision outcome – an area in which more research is needed (Fjermestad and Hiltz, 1999). Lastly, participants in the third study commented on the efficiencies gained by using an NKD tool in the idea-generation process. This is a promising finding and suggests that using an NKD tool in conjunction with a GSS for decision making may result in process gains. Past research has provided discouraging results regarding the process gains associated with GSS in comparison to face-to-face decision processes (Fjermestad and Hiltz, 1999). Thus, the results of this research are encouraging.

10.3 Strengths

In addition to the contributions to both research and practice noted above, this research has a number of other strengths. This research adopted a cross-level approach (Rousseau, 1985), examining knowledge discovery at the tool, individual, group and organizational levels. As discussed above, a competing-theories approach was used to provide a richer understanding of knowledge discovery.

Another strength is the diversity of the research methods employed. For example, both quantitative and qualitative research methods, as well as experiments and case study were used. A lab experiment with student participants was used to gain statistical power, whereas field studies were used to explore the phenomenon of novel-knowledge discovery in more depth. Further, both design-science and behavioural-science approaches were used.

The three studies conducted in this research were distinct, yet designed to support and triangulate to the results of the other studies. As discussed above, the iterative nature of the design-science process allowed each study to build on and validate the results of the previous studies.

10.4 Limitations

There are limitations to this research, which may affect the generalizability and strength of the results. The use of a single organizational setting in the case study and single group in the field experiment may limit the generalizability of the results. However, the different industries examined across these two field studies allows the results to be generalized to organizations that are interested in innovation, such as pharmaceutical and high-tech firms.

The prototypical state of the Athens tool was another limitation of this research. The Athens tool, not a commercial-grade tool, was compared to other commercial-grade tools such as Google and Vivisimo. Although the Athens training included comments regarding its state of development, the comparison to commercial-grade tools may have been unfair and limited the

validity of the results. However, when developing and evaluating new tools, this limitation is unavoidable.

The lengthy processing time of the Athens tool limited the degree to which participants could use the tool iteratively. Across all studies, the researcher completed the Athens queries prior to the start of the study. For the purposes of these studies, this limitation was unavoidable. However, the Athens tool is not designed to be used in the same iterative manner as tools such as Google and Vivisimo.

The other tool limitation is the combined novelty of both the tool and task. Across all research studies, participants were asked to use a highly novel tool to complete a highly novel task. The resulting cognitive complexity may have resulted in some participants reframing their task as a level 1 task and using the Athens tool according to their existing technological frames for Web-search tools.

The results of Study 1 – lab experiment – may have been limited by the use of student participants. Although the only feasible way to achieve the sample size required for sufficient statistical power, student participants may not have taken the task as seriously as organizational participants. Lack of organizational work experience may have limited the students' ability to complete the experimental task. However, all student participants were business students and, thus, have some familiarity with strategic tasks such as the one used in Study 1. Further, prior research suggests that the use of student subjects is appropriate (e.g. Greenberg et al., 1987).

10.5 Future Research

A number of future research opportunities were identified above. The results of this research provide several interesting avenues for future research. Although both group- and organizational-level learning were explored in this research, the emphasis was on examining individual-level learning. Thus, future research should more fully explore the group and organizational learning processes – interpreting, integrating and institutionalizing – associated

with novel-knowledge discovery (NKD). Given organizational resource constraints, more ideas are typically generated than can be implemented. Thus, it is important to understand how the ideas generated by individuals, with the support of an NKD tool, actually get implemented and institutionalized. Understanding this process can help organizations structure their processes, evaluation criteria and information systems for novel-knowledge discovery.

Other avenues for future research include exploring in more depth the impact that technological frames have on the use, design and effectiveness of NKD tools; exploring how an NKD tool is used and implemented in practice; examining where an NKD tool helps in the creativity and brainstorming process – once individual creativity has been exhausted, or before?; examining which design principles are most important for novel-knowledge discovery; testing the remaining cells of the external knowledge-discovery theoretical framework; exploring *internal* knowledge discovery in organizations and examining how the design, use and effectiveness of an NKD tool would change with an internal-focus.

10.6 Conclusions

This research examined novel-knowledge discovery in the context of organizational learning and external-knowledge discovery. The field studies conducted in this research helped validate the importance and challenges of novel-knowledge discovery to organizations. Despite the importance of novel knowledge to these organizations, there were no distinct processes or tools used specifically for its discovery. Novel-knowledge discovery was embedded in other processes and often the result of serendipity. An NKD design theory was developed to help address the challenges associated with novel-knowledge discovery and a tool instantiation and process of use evaluated. Despite the influence of existing technological frames for Web-search tools, the NKD tool – Athens – was perceived as a useful tool by organizational participants. In both field studies, participants indicated a desire to continue experimenting with the tool. In addition, participant comments about the efficiency of using the NKD tool Athens in combination

with decision tools such as GSS and collaborative filtering to generate and evaluate novel ideas provides promising evidence regarding the importance and usefulness of novel-knowledge discovery tools in organizations.

BIBLIOGRAPHY

- Abernathy, W. (1978) *The Productivity Dilemma*, Johns Hopkins University Press, Baltimore, MD.
- Agarwal, R. and Prasad, J. (1998) A conceptual and operational definition of personal innovativeness in the domain of information technology, *Information Systems Research*, 9, 2, 204-215.
- Agarwal, R., Sambamurthy, V., and Stair, R.M. (2000) Research report: The evolving relationship between general and specific computer self-efficacy - An empirical assessment, *Information Systems Research*, 11, 4, 418-430.
- Alavi, M. and Leidner, D.E. (2001) Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues, *MIS Quarterly*, 25, 1, 107-136.
- Allport, F. (1920) The influence of the group upon association and thought, *Journal of Experimental Psychology*, 3, 159-182.
- Amabile, T.M. (1983) The social psychology of creativity: A componential conceptualization, *Journal of Personality and Social Psychology*, 45, 2, 357-376.
- Amabile, T.M. (1988) A model of creativity and innovation in organizations, in B.M. Staw and L.L. Cummings (eds.) *Research in Organizational Behavior*, Jai Press Inc, Greenwich, Connecticut, 123-167.
- Amabile, T.M., Conti, R., Coon, H., Lazenby, J., and Herron, M. (1996) Assessing the work environment for creativity, *Academy of Management Journal*, 39, 5, 1154-1184.
- Anderson, J.R. (1987) Skill acquisition: Compilation of weak-method problem solutions, *Psychological Review*, 94, 2, 192-210.
- Anderson, J.R. (1992) Automaticity and the ACT* theory, *American Journal of Psychology*, 105, 2, 165-180.

- Anderson, P. and Tushman, M.L. (1990) Technological discontinuities and dominant designs: A cyclical model of technological change, *Administrative Science Quarterly*, 35, 4, 604-633.
- Angle, H.L. (1989) Psychology and organizational innovation, in A.H. Van de Ven, H.L. Angle and M.S. Poole (eds.) *Research on the management of innovation: The Minnesota studies*, Harper & Row, New York, 135-170.
- Anonymous (2005) W3C Semantic Web, <http://www.w3.org/2001/sw/> [Accessed June 13, 2005].
- Argote, L. (1999) *Organizational learning: Creating, retaining, and transferring knowledge*, Kluwer Academic Publishers, Boston.
- Argyris, C. and Schon, D.A. (1978) *Organizational learning*, Addison-Wesley, Reading, MA.
- Barnett, H.G. (1953) *Innovation: The Basis of Cultural Change*, McGraw-Hill Book Company, Inc., New York.
- Barron, F. and Harrington, D.M. (1981) Creativity, intelligence, and personality, *Annual Review of Psychology*, 32, 439-476.
- Barsalou, L.W. (1983) Ad hoc categories, *Memory and Cognition*, 11, 3, 211-227.
- Barsalou, L.W. (1992) Frames, concepts, and conceptual fields, in E. Kittay and A. Lehrer (eds.) *Frames, fields, and contrasts: New essays in semantic and lexical organization*, Erlbaum, Hillsdale, NJ, 21-74.
- Beccerra-Fernandez, I., Gonzalez, A., and Sabherwal, R. (2004) *Knowledge management: Challenges, solutions, and technologies*, Pearson Prentice Hall, Upper Saddle River, NJ.
- Benbasat, I. (1990) Laboratory experiments in information systems with a focus on individuals: A critical appraisal, in I. Benbasat (ed.) *The information systems research challenge: experimental research methods*, Harvard Business School, 33-47.
- Benbasat, I., Goldstein, D.K., and Mead, M. (1987) The case research strategy in studies of information-systems, *MIS Quarterly*, 11, 3, 369-386.

- Berners-Lee, T., Hendler, J., and Lassila, O. (2001) The Semantic Web, *Scientific American*, May 2001.
- Bierly, P. and Chakrabarti, A. (1996) Generic Knowledge Strategies in the U.S. Pharmaceutical Industry, *Strategic Management Journal*, 17, Winter Special Issue, 123-135.
- Blum, C. and Roli, A. (2003) Metaheuristics in combinatorial optimization: Overview and conceptual comparison, *ACM Computing Surveys*, 35, 3, 268-308.
- Boisot, M. (2002) The Creation and Sharing of Knowledge, in C.W. Choo and N. Bontis (eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge*, Oxford University Press, New York NY, 65-77.
- Boland, R.J. and Tenkasi, R.V. (1995) Perspective Making and Perspective-Taking in Communities of Knowing, *Organization Science*, 6, 4, 350-372.
- Boland, R.J., Tenkasi, R.V., and Teeni, D. (1994) Designing information technology to support distributed cognition, *Organization Science*, 5, 3, 456-475.
- Bontis, N., Crossan, M.M., and Hulland, J. (2002) Managing an organizational learning system by aligning stocks and flows, *Journal of Management Studies*, 39, 4, 437-469.
- Bowman, C.M., Danzig, P.B., Manber, U., and Schwartz, M.F. (1994) Scalable Internet resource discovery: research problems and approaches, *Communications of the ACM*, 37, 8, 98-107.
- Campbell, D.T. and Stanley, J.C. (1963) Experimental and quasi-experimental designs for research, Rand McNally, Chicago.
- Carlile, P.R. (2002) A pragmatic view of knowledge and boundaries: Boundary objects in new product development, *Organization Science*, 13, 4, 442-455.
- Carlile, P.R. (2004) Transferring, translating, transforming: An integrative framework for managing knowledge across boundaries, *Organization Science*, 15, 5, 555-568.
- Chen, H. "Knowledge Management Systems: A Text Mining Perspective," The University of Arizona, Tucson, AZ, 2001.

- Chen, H., Chau, M., and Zeng, D. (2002) CI Spider: A tool for competitive intelligence on the Web, *Decision Support Systems*, 34, 2002, 1-17.
- Choo, C.W. (2002) Sensemaking, Knowledge Creation, and Decision Making: Organizational Knowing as Emergent Strategy, in C.W. Choo and N. Bontis (eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge*, Oxford University Press, New York, 79-88.
- Chung, W., Chen, H., and Nunamaker Jr, J.F. (2005) A Visual Framework for Knowledge Discovery on the Web: An Empirical Study of Business Intelligence Exploration, *Journal of Management Information Systems*, 21, 4, 57-84.
- Churchill, G.A. (1979) A paradigm for developing better measures of marketing constructs, *Journal of Marketing Research*, 16, 1, 64-73.
- Cohen, J. (1960) A coefficient of agreement for nominal scales, *Educational and Psychological Measurement*, 20, 37-46.
- Cohen, J. (1988) *Statistical Power Analysis for the Behavioral Sciences*, Academic Press, New York.
- Cohen, W.M. and Levinthal, D.A. (1990) Absorptive capacity: A new perspective on learning and innovation, *Administrative Science Quarterly*, 35, 1, 128-152.
- Collins, H.M. (1993) The Structure of Knowledge, *Social Research*, 60, 1, pp. 95-116.
- Compeau, D.R. and Higgins, C.A. (1995) Application of social cognitive theory to training for computer skills, *ISR*, 6, 2, 118-143.
- Comrey, A.L. and Lee, H.B. (1992) *A First Course in Factor Analysis*, (2nd ed.), Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ.
- Costello, A.B. and Osborne, J.W. (2005) Best Practices in Exploratory Factor Analysis: Four Recommendations for Getting the Most From Your Analysis, *Practical Assessment Research & Evaluation*, 10, 7, 1-9.

- Couger, J.D. (1990) Ensuring Creative Approaches in Information System Design, *Managerial and Decision Economics*, 11, 5, 281-295.
- Couger, J.D. and Dengate, G. (1992) Measurement of Creativity in I.S. Products, in *Hawaii International Conference on Systems Sciences*, Lihue, HI, 288-298.
- Couger, J.D., Higgins, L.F., and McIntyre, S.C. (1993) (Un)structured creativity in information systems organizations, *MIS Quarterly*, 17, 4, 375-397.
- Crossan, M.M. and Berdrow, I. (2003) Organizational learning and strategic renewal, *Strategic Management Journal*, 24, 1087-1105.
- Crossan, M.M., Lane, H.W., and White, R.E. (1999) An organizational learning framework: From intuition to institution, *Academy of Management Review*, 24, 2, 522-537.
- Csikszentmihalyi, M. (1996) *Creativity: Flow and the Psychology of Discovery and Invention*, Harper Perennial, New York.
- Daft, R.L. and Weick, K.E. (1984) Toward a model of organizations as interpretation systems, *Academy of Management Review*, 9, 2, 284-295.
- Dahlin, K.B. and Behrens, D.M. (2005) When is an invention really radical? Defining and measuring technological radicalness, *Research Policy*, 34, 717-737.
- Davenport, T.H. and Volpel, S.C. (2001) The rise of knowledge towards attention management, *Journal of Knowledge Management*, 5, 3, 212-221.
- Day, G.S. (2002) Managing the market learning process, *The Journal of Business & Industrial Marketing*, 17, 4, 240-252.
- Dennis, A.R., Wixom, B.H., and Vandenberg, R.J. (2001) Understanding fit and appropriation effects in group support systems via meta-analysis, *MIS Quarterly*, 25, 2, 167-193.
- Desanctis, G. and Gallupe, R.B. (1987) A foundation for the study of group decision support systems, *Management Science*, 33, 5, 589-609.
- Desanctis, G. and Poole, M.S. (1994) Capturing the Complexity in Advanced Technology Use - Adaptive Structuration Theory, *Organization Science*, 5, 2, 121-147.

- Diehl, M. and Stroebe, W. (1987) Productivity loss in brainstorming groups: Toward the solution of a riddle, *Journal of Personality and Social Psychology*, 53, 3, 497-509.
- Diehl, M. and Stroebe, W. (1991) Productivity loss in idea-generating groups: Tracking down the blocking effects, *Journal of Personality and Social Psychology*, 61, 3, 392-403.
- Dietterich, T.G. (2000) An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization, *Machine Learning*, 40, 139-157.
- Dutta, D.K. and Crossan, M.M. (2005) The nature of entrepreneurial opportunities: Understanding the process using the 4I organizational learning framework, *Entrepreneurship Theory and Practice*, 29, 4, 425-449.
- Dworman, G.O., Kimbrough, S.O., and Patch, C. (2000) On pattern-directed search of archives and collections, *Journal of the American Society for Information Science*, 51, 1, 14-23.
- Edmondson, A.C. (2002) The local and variegated nature of learning in organizations: A group-level perspective, *Organization Science*, 13, 2, 128-146.
- Encinar, M.-I. and Muñoz, F.-F. (2006) On novelty and economics: Schumpeter's paradox, *Journal of Evolutionary Economics*, 16, 3, 255-277.
- Fabrigar, L.R., Wegener, D.T., MacCallum, R.C., and Strahan, E.J. (1999) Evaluating the Use of Exploratory Factor Analysis in Psychological Research, *Psychological Methods*, 4, 3, 272-299.
- Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P., and Uthurusamy, R. (1996) *Advances in Knowledge Discovery and Data Mining*, AAAI Press/The MIT Press, Menlo Park, CA.
- Fiol, C.M. and Lyles, M.A. (1985) Organizational learning, *Academy of Management Review*, 10, 4, 803-813.
- Fjermestad, J. and Hiltz, S.R. (1999) An assessment of group support systems experimental research: Methodology and results, *Journal of Management Information Systems*, 15, 3, 7-149.

- Friendly, M. (2006) Power Analysis Help, *Retrieved from the World Wide Web on May 26, 2006 from <http://www.math.yorku.ca/SCS/Online/power/help.html#effect>*.
- Gallupe, R.B. and McKeen, J.D. (1990) Enhancing Computer-Mediated Communication - an Experimental Investigation into the Use of a Group Decision Support System for Face-to-Face Versus Remote Meetings, *Information & Management*, 18, 1, 1-13.
- Ghiselin, B.E. (1952) *The Creative Process*, University of California Press, Berkeley.
- Gnyawali, D.R. and Stewart, A.C. (2003) A contingency perspective on organizational learning: Integrating environmental context, organizational learning processes, and types of learning, *Management Learning*, 34, 1, 63-89.
- Gogan, J.L. (2006) Commentary on Karl E. Weick's 'The role of imagination in the organizing of knowledge', *European Journal of Information Systems*, 15, 453-456.
- Gordon, M.D. and Lindsay, R.K. (1996) Toward discovery support systems: A replication, re-examination, and extension of Swanson's work on literature-based discovery of a connection between Raynaud's and fish oil, *Journal of the American Society for Information Science*, 47, 2, 116-128.
- Grant, R.M. (1996) Toward a knowledge-based theory of the firm, *Strategic Management Journal*, 17, 109-122.
- Gray, P.H. and Meister, D.B. (2004) Knowledge Sourcing Effectiveness, *Management Science*, 50, 6, 821-834.
- Greenberg, J., Gordon, M.E., Slade, L.A., and Schmitt, N. (1987) The College Sophomore as Guinea Pig: Setting the Record Straight/Student Guinea Pigs: Porcine Predictors and Particularistic Phenomena, *Academy of Management. The Academy of Management Review*, 12, 1, 157-159.
- Gregor, S. (2006) The nature of theory in information systems, *MIS Quarterly*, 30, 3, 611-642.
- Grobelnik, M. and Mladenic, D. (2005) Automated knowledge discovery in advanced knowledge management, *Journal of Knowledge Management*, 9, 5, 132-149.

- Guilford, J.P. (1956) The structure of intellect, *Psychological Bulletin*, 53, 4, 267-293.
- Hall, J.K. and Martin, M.J.C. (2005) Disruptive technologies, stakeholders and the innovation value-added chain: a framework for evaluating radical technology development, *R & D Management*, 35, 3, 273-284.
- Hargadon, A.B. (2002) Brokering knowledge: Linking learning and innovation, *Research in Organizational Behaviour*, 24, 41-85.
- Hargadon, A.B. and Douglas, Y. (2001) When innovations meet institutions: Edison and the design of the electric light, *Administrative Science Quarterly*, 46, 476-501.
- Hargadon, A.B. and Fanelli, A. (2002) Action and possibility: Reconciling dual perspectives of knowledge in organizations, *Organization Science*, 13, 3, 290-302.
- Hargadon, A.B. and Sutton, R.I. (1997) Technology brokering and innovation in a product development firm, *Administrative Science Quarterly*, 42, 716-749.
- Harrison, J.S. and St. John, C.H. (2004) Foundations in Strategic Management, (3rd ed.), South-Western, Thomson Corporation, Mason, OH.
- Hedberg, B. (1981) How organizations learn and unlearn?, in P.C. Nystrom and W.H. Starbuck (eds.) *Handbook of organizational design*, Oxford University Press, London.
- Hender, J.M., Dean, D.L., Rodgers, T.L., and Nunamaker Jr., J.F. (2002) An Examination of the Impact of Stimuli Type and GSS Structure on Creativity: Brainstorming Versus Non-Brainstorming Techniques in a GSS Environment., *Journal of Management Information Systems*, 18, 4, 59-85.
- Hevner, A.R., March, S.T., Park, J., and Ram, S. (2004) Design science in Information Systems research, *MIS Quarterly*, 28, 1, 75-105.
- Hill, C.W.L. and Rothaermel, F.T. (2003) The performance of incumbent firms in the face of radical technological innovation, *Academy of Management. The Academy of Management Review*, 28, 2, 257-274.

- Hintzman, D.L. (1988) Judgements of frequency and recognition memory in a multiple-trace memory model, *Psychological Review*, 95, 4, 528-551.
- Hoare, C. and Sorensen, H. (2005) Information Foraging with a Proximity-Based Browsing Tool, *The Artificial Intelligence Review*, 24, 3-4, 233-252.
- Hoffman, L.R. (1959) Homogeneity of member personality and its effect on group problem solving, *Journal of Abnormal and Social Psychology*, 58, 27-32.
- Hsu, M.-H. and Chiu, C.-M. (2004) Internet self-efficacy and electronic service acceptance, *Decision Support Systems*, 38, 369-381.
- Huber, G.P. (1991) Organizational learning: The contributing processes and the literatures, *Organization Science*, 2, 1, Special Issue: Organizational Learning: Papers in Honor of (and by) James G. March, 88-115.
- Jarvis, C.B., Mackenzie, S.B., and Podsakoff, P.M. (2003) A critical review of construct indicators and measurement model misspecification in marketing and consumer research, *Journal of Consumer Research*, 30, 199-218.
- Jick, T.D. (1979) Mixing qualitative and quantitative methods: Triangulation in action, *Administrative Science Quarterly*, 24, 4, 602-611.
- Johnson, R.D. and Marakas, G.M. (2000) Research report: The role of behavioural modeling in computer skills acquisition - toward a refinement of the model, *Information Systems Research*, 11, 4, 402-417.
- Kane, G.C. and Alavi, M. (2005) Information technology and organizational learning: An investigation of exploitation and exploration processes, in *Proceedings of the Twenty-Sixth International Conference on Information Systems*, Las Vegas 459-470.
- Kane, G.C. and Alavi, M. (2007) Information technology and organizational learning: An investigation of exploration and exploitation processes, *Organization Science*, 18, 5, 796-812.

- Kanter, R.M. (1988) When a thousand flowers bloom: Structural, collective and social conditions for innovation in organization, in B.M. Staw and L.L. Cummings (eds.) *Research in Organizational Behavior*, Jai Press Inc, Greenwich, Connecticut, 169-211.
- Kim, D.H. (1993) The link between individual and organizational learning, *Sloan Management Review*, 35, 1, 37-50.
- Kimbrough, S.O. (2001) Core of Discovery Executive Briefing: Capabilities of Practical Reasoning's Core of Discovery, <http://opim-sun.wharton.upenn.edu/~sok/asadai/cod-exec-brief.html>. [Accessed Jan 7, 2007].
- Kirton, M. (1976) Adaptors and innovators: A description and measure, *Journal of Applied Psychology*, 61, 5, 622-629.
- Kirton, M.J. (2003) *Adaption-Innovation: In the Context of Diversity and Change*, Routledge, New York.
- Knowles, M.S. (1970) *The Modern Practice of Adult Education: Andragogy versus Pedagogy*, Association Press, New York.
- Kourteli, L. (2005) Scanning the business external environment for information: evidence from Greece, *Information Research*, 11, 1, 1-16.
- Lamm, H. and Trommsdorff, G. (1973) Group versus individual performance on tasks requiring ideational proficiency (brainstorming), *European Journal of Social Psychology*, 3, 4, 361-388.
- Lana, R.C. (1969) Pretest sensitization, in R. Rosenthal and R.L. Rosnow (eds.) *Artifact in behavioral research*, Academic Press, New York, 119-141.
- Leidner, D.E. and Jarvenpaa, S.L. (1995) The use of information technology to enhance management school education: A theoretical view, *MIS Quarterly*, 19, 3, 265-291.
- Leifer, R., McDermott, C.M., O'Connor, G.C., Peters, L.S., Rice, M.P., Veryzer, R.W., and Rice, M. (2000) *Radical Innovation*, Harvard Business School Press, Boston, MA.

- Levitt, B. and March, J.G. (1988) Organizational learning, *Annual Review of Sociology*, 14, 319-340.
- Lewis, W., Agarwal, R., and Sambamurthy, V. (2003) Sources of influence on beliefs on information technology use: An empirical study of knowledge workers, *MIS Quarterly*, 27, 4, 657-678.
- Lytras, M. (2005) Semantic Web and Information Systems: An Agenda Based on Discourse with Community Leaders, *International Journal of Semantic Web and Information Systems*, 1, 1, i-xii.
- MacCallum, R.C., Widaman, K.F., Zhang, S., and Hong, S. (1999) Sample Size in Factor Analysis, *Psychological Methods*, 4, 1, 84-99.
- Mackay, J.M. and Elam, J.J. (1992) A comparative study of how experts and novices use a decision aid to solve problems in complex knowledge domains, *Information Systems Research*, 3, 2, 150-172.
- Maier, N.R.F. (1945) Reasoning in humans: III. The mechanisms of equivalent stimuli and of reasoning, *Journal of Experimental Psychology*, 35, 349-360.
- Majchrzak, A., Cooper, L.P., and Neece, O.E. (2004) Knowledge reuse for innovation, *Management Science*, 50, 2, 174-188.
- Majchrzak, A., Malhotra, A., and John, R. (2005) Perceived individual collaboration know-how development through information technology-enabled contextualization: Evidence from distributed teams, *Information Systems Research*, 16, 1, 9-27.
- Marakas, G.M., Yi, M.Y., and Johnson, R.D. (1998) The multilevel and multifaceted character of computer self-efficacy: Toward clarification of the construct and an integrative framework for research, *Information Systems Research*, 9, 2, 126-163.
- March, J.G. (1991) Exploration and exploitation in organizational learning, *Organization Science*, 2, 1, Special Issue: Organizational Learning: Papers in Honour of (and by) James G. March, 71-87.

- March, S.T. and Smith, G.F. (1995) Design and Natural-Science Research on Information Technology, *Decision Support Systems*, 15, 4, 251-266.
- Markus, M.L., Majchrzak, A., and Gasser, L. (2002) A design theory for systems that support emergent knowledge processes, *MIS Quarterly*, 26, 3, 179-212.
- Mason, R.O. (1990) MIS Experiments: A Pragmatic Perspective, in I. Benbasat (ed.) *The Information Systems Research Challenge: Experimental Research Methods*, Harvard Business School, 3-20.
- McGrath, J.E. (1984) *Groups: Interaction and Performance*, Prentice Hall, Englewood Cliffs, NJ.
- McGrath, J.E. and Altman, I. (1966) *Small Group Research: A Synthesis and Critique of the Field*, Holt, Rinehart and Winston, New York.
- McLean, L.D. (2005) Organizational culture's influence on creativity and innovation: A review of the literature and implications for human resource development, *Advances in Developing Human Resources*, 7, 2, 226-246.
- Mednick, S. (1962) The associative basis of the creative process, *Psychological Review*, 29, 220-232.
- Miles, M.B. and Huberman, M.A. (1994) *Qualitative Data Analysis: An Expanded Sourcebook*, Sage Publications, Inc., Thousand Oaks.
- Mooney, R.L. (1963) A Conceptual Model for Integrating Four Approaches to the Identification of Creative Talent, in C.W. Taylor and F. Barron (eds.) *Scientific Creativity: Its Recognition and Development*, John Wiley & Sons, Inc., New York.
- Moore, G.C. and Benbasat, I. (1991) Development of an Instrument to Measure the Perceptions of Adopting an Information Technology Innovation, *Information Systems Research*, 2, 3, 192-222.
- Mumford, M.D. and Gustafson, S.B. (1988) Creativity syndrome: Integration, application and innovation, *Psychological Bulletin*, 103, 1, 27-43.

- Nagasundaram, M. and Bostrom, R.P. (1995) The structuring of creative processes using GSS: A framework for research, *Journal of Management Information Systems*, 11, 3, 87-114.
- Nagasundaram, M. and Dennis, A.R. (1993) When a group is not a group: The cognitive foundations of group idea generation, *Small Group Research*, 24, 4, 463-489.
- Nemeth, C.J. (1986) Differential contributions of majority and minority influence, *Psychological Review*, 93, 1, 23-32.
- Newell, A., Shaw, J.C., and Simon, H.A. (1962) The process of creative thinking, in H.E. Gruber, G. Terrell and M. Wertheimer (eds.) *Contemporary Approaches to Creative Thinking*, Atherton Press, New York.
- Nonaka, I. (1994) A dynamic theory of organizational knowledge creation, *Organization Science*, 5, 1, 14-37.
- Norman, D.A. (1982) *Learning and Memory*, W.H. Freeman, San Francisco.
- Nunnally, J.C. (1978) *Psychometric Theory*, McGraw-Hill, New York.
- Nystrom, P.C. and Starbuck, W.H. (1984) To avoid organizational crises, unlearn, *Organizational Dynamics*, 12, 4, 53-65.
- Ocker, R., Hiltz, S.R., Turoff, M., and Fjermestad, J. (1996) The effects of distributed group support and process structuring on software requirements development teams: Results on creativity and quality, *Journal of Management Information Systems*, 12, 3, 127-153.
- Orlikowski, W.J. (2002) Knowing in practice: Enacting a collective capability in distributed organizing, *Organization Science*, 13, 3, 249-273.
- Orlikowski, W.J. and Barley, S.R. (2001) Technology and institutions: What can research on information technology and research on organizations learn from each other?, *MIS Quarterly*, 25, 2, 145-165.
- Orlikowski, W.J. and Gash, D.C. (1991) Changing frames: Towards an understanding of information technology and organizational change, in *Academy of Management Best Paper Proceedings*, 51st Annual Meeting, Miami Beach, FL, August, 189-193.

- Orlikowski, W.J. and Gash, D.C. (1994) Technological frames: Making sense of information technology in organizations, *ACM Transactions on Information Systems*, 12, 2, 174-207.
- Orlikowski, W.J. and Iacono, C.S. (2001) Research commentary: Desperately seeking the "IT" in IT research - A call to theorizing the IT artifact, *Information Systems Research*, 12, 2, 121-134.
- Osborn, A.F. (1953) *Applied Imagination*, Charles Scribner's Sons, New York.
- Parker, M.M., Benson, R.J., and Trainor, H.E. (1988) *Information Economics. Linking Business Performance to Information Technology*, Prentice Hall, Englewood Cliffs, New Jersey.
- Pearson, P.H. (1970) Relationships between global and specified measures of novelty-seeking, *Journal of Consulting and Clinical Psychology*, 34, 199-204.
- Penrose, E.T. (1959) *The theory of the growth of the firm*, (U.S. ed.), M. E. Sharpe, White Plains, NY.
- Peters, M. and Robinson, V. (1984) The Origins and Status of Action Research, *The Journal of Applied Behavioral Science*, 20, 2, 113-124.
- Piaget, J. (1954) *The Construction of Reality in the Child*, Basic Books, New York.
- Piaget, J. and Inhelder, B. (1969) *The Psychology of the Child*, Basic Books, Inc., New York.
- Piccoli, G., Ahmad, R., and Ives, B. (2001) Web-based virtual learning environments: A research framework and a preliminary assessment of effectiveness in basic IT skills training, *MIS Quarterly*, 25, 4, 401-426.
- Pirolli, P. and Card, S. (1995) Information foraging in information access environments, in *Proceedings of the SIGCHI conference on Human factors in computing systems*, Denver, Colorado, United States, ACM Press/Addison-Wesley Publishing Co., 51-58.
- Pirolli, P. and Card, S. (1999) Information foraging, *Psychological Review*, 106, 4, 643-675.
- Polanyi, M. (1966) *The Tacit Dimension*, Doubleday, Garden City, N. Y.
- Potter, R.E. and Balthazard, P. (2004) The role of individual memory and attention processes during electronic brainstorming, *MIS Quarterly*, 28, 4, 621-643.

- Quinn, J.B. (1996) Team innovation, *Executive Excellence*, 13, 7, 13-14.
- Rao, R. (2004) Leveraging content in enterprise knowledge processes, in M. Rao (ed.) *Knowledge scaffolding: How practitioners use knowledge management tools*, Butterworth-Heinemann.
- Robey, D. and Boudreau, M.C. (1999) Accounting for the contradictory organizational consequences of information technology: Theoretical directions and methodological implications, *Information Systems Research*, 10, 2, 167-185.
- Rousseau, D.M. (1985) Issues of Level in Organizational Research: Multi-level and Cross-level Perspectives, in L.L. Cummings and B.M. Staw (eds.) *Research in Organizational Behavior*, JAI Press Inc., Greenwich, CT, 1-37.
- Ryan, S.D. and Harrison, D.A. (2000) Considering social subsystem costs and benefits in information technology investment decisions: A view from the field on anticipated payoffs, *Journal of Management Information Systems*, 16, 4, 11-40.
- Sambamurthy, V., Bharadwaj, A., and Grover, V. (2003) Shaping agility through digital options: Reconceptualizing the role of information technology in contemporary firms, *MIS Quarterly*, 27, 2, 237-263.
- Sambamurthy, V. and Poole, M.S. (1992) The effects of variations in capabilities of GDSS designs on management of cognitive conflict in groups, *Information Systems Research*, 3, 3, 224-251.
- Scandura, T.A. and Williams, E.A. (2000) Research Methodology in Management: Current Practices, Trends, and Implications for Future Research, *Academy of Management Journal*, 43, 6, 1248-1264.
- Schulz, M. (2001) The uncertain relevance of newness: Organizational learning and knowledge flows, *Academy of Management Journal*, 44, 4, 661-681.

- Schumpeter, J.A. ([1932] 2005) *Development*. Translated by Becker, M.C. and Knudsen, T. With an introduction by Becker, M.C., EBlinger, H.U., Hedtka, U. and Knudsen, T., *Journal of Economic Literature*, XLIII, 1, 108–120.
- Schweizer, T.S. (2006) The Psychology of Novelty-Seeking, Creativity and Innovation: Neurocognitive Aspects Within a Work-Psychological Perspective, *Creativity and Innovation Management*, 15, 2, 164-172.
- Scotchmer, S. and Green, J. (1990) Novelty and Disclosure in Patent Law, *The Rand Journal of Economics*, 21, 1, 131-146.
- Shadish, W.R., Cook, T.D., and Campbell, D.T. (2003) *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*, Houghton Mifflin Company, Boston, MA.
- Shaw, M.E. (1973) Scaling group tasks: A method for dimensional analysis, *JSAS Catalog of Selected Documents in Psychology*, 3, 8.
- Shaw, M.E. (1976) *Group Dynamics: The Psychology of Small Group Behavior*, McGraw Hill, Inc., New York.
- Shneiderman, B. (1996) The eyes have it: A task by data type taxonomy for information visualizations, in *IEEE Symposium on Visual Languages*, Los Alamitos, CA, IEEE Computer Society Press, 339-343.
- Shneiderman, B. (2000) Creating creativity: User interfaces for supporting innovation, *ACM Transactions on Computer-Human Interaction*, 7, 1, 114-138.
- Simon, H.A. (1969) *Sciences of the Artificial*, M.I.T. Press, Cambridge, MA.
- Skillicorn, D.B. "Best global assessments from individual opinions," in: *Unpublished manuscript*, 2001, pp. 1-5.
- Skillicorn, D.B. (2007) *Understanding Complex Datasets: Data Mining with Matrix Decompositions*, CRC Press, Boca Raton, FL.

- Skillicorn, D.B. and Vats, N. (2004) Novel Information Discovery for Intelligence and Counterterrorism, Department of Computing and Information Science, Queen's University, External Technical Report 2004-488.
- Stein, E.W. and Zwass, V. (1995) Actualizing organizational memory with information systems, *Information Systems Research*, 6, 2, 85-117.
- Steiner, I.D. (1972) Group Process and Productivity, Academic Press, New York.
- Stone, E.F. (1978) Research Methods in Organizational Behaviour, Goodyear Publishing Co., Santa Monica.
- Stoner, J.A.F. (1968) Risky and cautious shifts in group decisions: The influence of widely held values, *Journal of Experimental Psychology*, 4, 442-459.
- Straub, D.W. (1989) Validating Instruments in MIS Research, *MIS Quarterly*, 13, 2, 147-169.
- Sturdivant, R. (2006) Power and Sample Size in ANOVA, Retrieved from the World Wide Web on May 25, 2006 from <http://www.dean.usma.edu/math/people/Sturdivant/A376%20lessons/Block%20one/376lsn21.htm>.
- Styhre, A. (2006) Organization Creativity and the Empiricist Image of Novelty, *Creativity and Innovation Management*, 15, 2, 143-149.
- Susman, G.I. and Evered, R.D. (1978) An Assessment of the Scientific Merits of Action Research, *Administrative Science Quarterly*, 23, 4, 582-603.
- Swanson, D.R. (1986) Fish oil, Raynaud's syndrome, and undiscovered public knowledge, *Perspectives in Biology and Medicine*, 30, 1, 7-18.
- Swanson, D.R. (1990) Medical literature as a potential source of new knowledge, *Bulletin of the Medical Library Association*, 78, 1, 29-37.
- Zsulanski, G. (1996) Exploring internal stickiness: impediments to the transfer of best practice within the firm, *Strategic Management Journal*, 17, Winter Special Issue, 27-43.
- Tabachnick, B.G. and Fidell, L.S. (2001) Using Multivariate Statistics, (4th ed.), Allyn & Bacon, Needham Heights, MA.

- Tan, P., Steinbach, M., and Kumar, V. (2006) Introduction to Data Mining, Addison-Wesley.
- Tillquist, J., King, J.L., and Woo, C. (2002) A representational scheme for analyzing information technology and organizational dependency, *MIS Quarterly*, 26, 2, 91-118.
- Tushman, M.L. and Anderson, P. (1986) Technological discontinuities and organizational environments, *Administrative Science Quarterly*, 31, 3, 439-465.
- Tushman, M.L. and O'Reilly III, C.A. (1999) Building ambidextrous organizations: Forming your own "skunk works", *Health Forum Journal*, 42, 2, 20-23.
- Vandenbosch, B. and Higgins, C. (1996) Information acquisition and mental models: An investigation into the relationship between behaviour and learning, *Information Systems Research*, 7, 2, 198-214.
- Vandenbosch, B. and Huff, S.L. (1997) Searching and scanning: How executives obtain information from executive information systems, *MIS Quarterly*, 21, 1, 81-107.
- VanGundy, A.B. (1988) Techniques of Structured Problem Solving, (2nd ed.), Van Nostrand Reinhold, New York.
- Vats, N. and Skillicorn, D.B. (2004a) The ATHENS system for novel information discovery, Department of Computing and Information Science, Queen's University, Technical Report 2004-489.
- Vats, N. and Skillicorn, D.B. (2004b) Information discovery within organizations using the Athens system, in *Proceedings of the 2004 Conference of the Centre for Advanced Studies on Collaborative Research*, Markham, Ontario 282-292.
- Venable, J. (2006) The role of theory and theorizing in design science research, in *First International Conference on Design Science Research in Information Systems and Technology*, February 24-25, Claremont, CA.
- Vinokur, A. (1969) Distribution of initial risk levels and group decisions involving risk, *Journal of Personality and Social Psychology*, 13, 3, 207-214.

- Vinokur, A. and Burnstein, E. (1974) Effects of partially shared persuasive arguments on group-induced shifts: A group problem-solving approach, *Journal of Personality and Social Psychology*, 29, 3, 305-315.
- Visser, E.-J. and Boschma, R. (2004) Learning in Districts: Novelty and Lock-in in a Regional Context, *European Planning Studies*, 12, 6, 793-808.
- Wallas, G. (ed.) *The Art of Thought*. J. Cape, London, 1926.
- Walls, J.G., Widmeyer, G.R., and El Sawy, O.A. (1992) Building an information system design theory for vigilant EIS, *Information Systems Research*, 3, 1, 36-59.
- Wang, C.L. and Ahmed, P.K. (2003) Organisational learning: a critical review, *The Learning Organization*, 10, 1, 8-17.
- Watson, R.T., DeSanctis, G., and Poole, M.S. (1988) Using a GDSS to facilitate group consensus: Some intended and unintended consequences, *MIS Quarterly*, 12, 3, 463-478.
- Watzlawick, P., Weakland, J.H., and Fisch, R. (1974) *Change*, Norton, New York.
- Webster, J. and Watson, R.T. (2002) Analyzing the past to prepare for the future: Writing a literature review, *MIS Quarterly*, 26, 2, xiii-xxiii.
- Weick, K.E. (1990) Technology as Equivoque: Sensemaking in New Technologies, in P.S. Goodman, L.S. Sproull and Associates (eds.) *Technology and Organizations*, Jossey-Bass Inc., San Francisco, CA.
- Weick, K.E. (2006) The role of imagination in the organizing of knowledge, *European Journal of Information Systems*, 15, 446-452.
- Weick, K.E. and Roberts, K.H. (1993) Collective mind in organizations: Heedful interrelating on flight decks, *Administrative Science Quarterly*, 38, 3, 357-381.
- West, M.A. and Anderson, N.R. (1996) Innovation in top management teams, *Journal of Applied Psychology*, 81, 6, 680-693.
- Witt, U. (2003) Evolutionary economics and the extension of evolution to the economy, in U. Witt (ed.) *The evolving economy*, Edward Elgar, Cheltenham UK.

- Yin, R.K. (1994) *Case study research: Design and methods*, Sage Publications, Thousand Oaks, CA.
- Zack, M.H. (1999) Developing a Knowledge Strategy, *California Management Review*, 41, 3, 125-145.
- Zack, M.H. (2001) If Knowledge is the Solution, then What's the Problem?, in Y. Malhotra (ed.) *Knowledge Management and Business Model Innovation*, Idea Group Publishing, Hershey, PA.
- Zack, M.H. (2002) Developing a knowledge strategy: Epilogue, in N. Bontis and C.W. Choo (eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge: A Collection of Readings*, Oxford University Press.
- Zack, M.H. (2005) The strategic advantage of knowledge and learning, *International Journal of Intellectual Capital and Learning*, 2, 1, 1-20.
- Zaheer, A. and Zaheer, S. (1997) Catching the wave: Alertness, responsiveness, and market influence in global electronic networks, *Management Science*, 43, 11, 1493.
- Zaltman, G., Duncan, R., and Holbeck, J. (1973) *Innovations and organizations*, Wiley, Chichester, England.
- Zigurs, I. and Buckland, B.K. (1998) A theory of task/technology fit and group support systems effectiveness, *MIS Quarterly*, 22, 3, 313-334.

APPENDIX A

SAMPLE SIZE

A.1 Sample Size Calculations

Determining sample size is an important aspect of research design because it affects the power of the study – the probability of finding a significant result. Calculating sample size is, in part, dependent on the type of statistical analysis being used. For MAN(C)OVA, a sample size of approximately 20 participants per cell is advised (Tabachnick and Fidell, 2001). Further, determining the appropriate sample size is dependent on the following parameters (Sturdivant, 2006):

1. Desired significance level.
2. Desired power level.
3. Estimated effect size.

To calculate the sample size for this study, I used a power level of 80%, as recommended by Cohen (1988) and a significance level of .05. The effect size represents the difference between group means in units of the within-group standard deviation (Friendly, 2006). In social science studies, effect sizes range from 0 to 3 (Cohen, 1988).

Statistical applications typically provide tools to help calculate sample size based on the criteria shown above. The statistical application R allows one to specify the number of groups, desired power level, desired significance level as well as the between and within-group variances (for effect size) and calculates the recommended sample size based on this input.

A.1.1 Sample Size for Lab Experiment

Based on the use of a 7-point Likert scale to measure the treatment group outcomes, I estimated a between-group variance of 0.33¹ and a within-group variance of 2 for the treatment

¹ R allows you to determine the between-group variance by specifying the estimated group means. For treatment conditions 1, 2 and 3, I used estimated means of 4, 4, 5 to calculate the between-group variances.

conditions in the lab experiment. Based on using 3 treatment groups, significance level of .05 and power level of 80%, the calculated sample size is 30. Thus, a minimum sample size of 30 participants per condition, for a total of 90 student participants, was used for experiment 1.

A.1.2 Sample Size for Field Experiment

Since organizational participants were used for the field experiment, it was infeasible to recruit enough participants to form the required number of groups to reach statistical significance. Instead, one group was used, composed of enough participants to form a realistic organizational group for a decision making task. According to Erik Lockhart, Associate Director of the Executive Decision Centre at Queen's School of Business, a realistic group is made up of between 6 to 12 participants. Thus, the goal was to recruit 8 to 10 participants from the local organization chosen to participate in this experiment.

APPENDIX B TREATMENT CONDITIONS

Table B.1 – Summary of Treatment Conditions and Experimental Task

	Treatment Condition Description	Tool	Process	Tests
1	Level 1 knowledge-discovery tool used by individuals.	Google	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2a, P2d, P3a, P4a, P4b, P4c, P5
2	Level 2 knowledge-discovery tool used by individuals.	Vivisimo	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2b, P2e, P3b, P4a, P4b, P4c, P5
3	Level 3 knowledge-discovery tool used by individuals.	Athens 2.0	Each individual used the tool to find, evaluate and suggest best idea to senior management.	P1, P2c, P2f, P3c, P4a, P4b, P4c, P5, P6a, P6b, P6c, P7a, P7b
4	Level 3 knowledge-discovery tool used by individuals in conjunction with collaborative-filtering tool.	Athens 2.0	Results from each individual's search were collected, combined and rated by each individual. Ratings were fed into a collaborative-filtering tool that recommended the best idea to senior management.	P6a, P6b, P6c, P7a, P7b
5	Level 3 knowledge-discovery tool used by individuals in conjunction with GSS tool.	Athens 2.0	Results from each individual's search were collected, combined and rated by each individual. Group collectively evaluated and suggested the best idea to senior management using GSS support structures to build consensus.	P6a, P6b, P6c, P7a, P7b

Experimental Task for Treatment Conditions 1, 2, 3 (Student participants)

1. Prior to the start of the experiment, participants were provided with a questionnaire in which pretest questions were asked. Participants were provided with the knowledge strategy question to address, determined prior to the start of the experiment. Each participant was free to interpret the question in his/her own manner, and develop search keywords. These search keywords were entered into the questionnaire.
2. At the start of the experiment, after all training was completed, each participant used the knowledge-discovery tool assigned to that condition to discover knowledge.
 - a. Participants using Google and Vivisimo were allowed to conduct multiple searches, since this is how individuals naturally use search tools, in order to generate what they viewed were the best ideas to address the knowledge strategy question. Each participant was asked to keep track of the ideas they generated (top 5 ideas, in their opinion), which they thought would potentially address the knowledge strategy question. They used a Web survey tool to capture their ideas and the associated Web link. There was a time limit of 45 minutes for the search task.
 - b. Participants using Athens were provided with the Athens results, which were generated by the experimenter the day before based on the search terms provided in the participant's questionnaire. The Athens tool requires 6 to 10 hours of processing time and therefore the "natural" use of the tool is not as iterative as the other search tools. Iteration is possible, but for the purposes of this study, participants did not use Athens iteratively. This is a more conservative use of Athens. Each participant was asked to keep track of the ideas they generated (top 5 ideas, in their opinion), which they thought would potentially address the knowledge strategy question. They used a Web survey tool to capture their ideas and the associated Web link. There was a time limit of 45 minutes to review the results and generate ideas.
3. Each participant evaluated his/her ideas and identified the best idea (to address the knowledge strategy question). The participant provided the associated links and described their best idea using one paragraph, which was then submitted to three independent evaluators.
4. At the end of the experiment, three evaluators were asked to separately evaluate the potential benefit and the radicalness of the chosen ideas on a 7-point Likert scale.
5. In addition, each participant rated each idea generated according to the potential benefit and radicalness on a 7-point Likert scale. Each participant also rank ordered the ideas by how well the idea addressed the task.
6. After the experiment, each individual was asked to answer questions about their experience with the tool in order to assess the type of knowledge discovered (deep, broad, novel), how the tool changed their understanding of the subject (mental model maintenance, tuning or building) and how the tool helped them learn (information foraging, search term development, intuiting, interpreting). Other potential confounds were measured at this time.

Experimental Task for Treatment Condition 3, 4, 5 (Organizational participants)

1. Training on Athens was conducted, including a test run of Athens to show participants the kinds of results returned by Athens. The training task was different from the knowledge strategy task used in the actual experiment.
2. After training was completed, participants were provided with a “pretest” questionnaire in which pretest questions were asked. In addition, participants were provided with a knowledge strategy question to address, determined prior to the start of the experiment.
3. Each individual was free to interpret the question in his/her own manner, and develop search keywords. Participants entered their search terms in the pretest questionnaire. The experimenter used these terms to run on Athens (multiple copies of Athens were running across multiple computers in the lab).
7. On the day of the experiment, the results from each individual’s search were displayed on the individual’s GSS display (each individual was assigned a specific seat/computer in the lab), but not yet shared with the group. The individual reviewed and evaluated the results, in search of the best idea to address the knowledge strategy question. Time limit: 45 minutes.
8. Using the GSS tool as an input mechanism, the individual rated the ideas generated by potential benefit, radicalness, and how well the idea addressed the task on a 7-point Likert scale.
4. The individual chose the best idea (to address the knowledge strategy question) from their list of ideas and described this idea in one paragraph, including a link to the Web page that spurred them to generate this idea. **(Treatment Condition 3)**
5. At this point in the experiment, each individual was asked to answer questions about their experience with the tool in order to assess the type of knowledge discovered (deep, broad, novel), how the tool changed their understanding of the subject (mental model maintenance, tuning or building) and how the tool helped them learn (information foraging, search term development, intuiting, interpreting). Satisfaction with the decision process and confidence were measured. Other potential confounds were measured at this time.
6. The ideas generated by each individual search were collected and combined into one list by the GSS facilitator using the GSS tool. The ideas were anonymous and, therefore, could not be linked to the individual who generated the idea.
7. The GSS facilitator, using the GSS tool, merged duplicate ideas in order to reduce the number of results to a manageable set.
8. Individuals reviewed and evaluated the reduced result set, in search of the best idea to address the knowledge strategy question. Time limit: 45 minutes.
9. The GSS facilitator, using the GSS tool, facilitated the voting process (across all ideas). During the voting process, each individual in the group rated the ideas generated by potential benefit, radicalness, and how well the idea addressed the task. These ratings were subsequently fed into a collaborative-filtering tool, which recommended the best idea. In addition, satisfaction with the decision process and confidence were measured via questionnaire. **(Treatment Condition 4)**
10. The GSS facilitator, using the GSS tool, facilitated a group discussion in which the group reviewed the voting results using a decision model and discussed the relevancy and importance of each idea.
11. After this group discussion, the individuals in the group were asked to vote on all the ideas a second time. The GSS facilitator facilitated another group discussion in which the group reviewed the voting results using a decision model and chose the best idea to address the knowledge strategy question. The group collectively described the idea in one paragraph. Satisfaction with the decision process and confidence were measured via questionnaire **(Treatment Condition 5)**.

APPENDIX C

THREATS TO VALIDITY

The following tables provide a detailed analysis of the threats to validity for the lab and field experiments, based on the types of threats identified in Shadish et al. (2003).

Table C.1 – Threats to Validity – Summary

Threat	Mitigation Strategy
Effect caused or impeded by the evaluator's opinion and biases.	Utilized five evaluators and checked for inter-rater reliability.
Effect caused or impeded by the GSS facilitator.	Tested the effect on individuals (no facilitator) separately from the group (facilitator).
Effect caused or impeded by the GSS.	Tested the effect on individuals (GSS functions not used) separately from the group (GSS functions used).
Effect caused or impeded by lack of training.	Prior to each experimental session, provided training to participants on the tool to be used. Asked individuals at the end of the experiment if they felt that training was sufficient to complete the task.
Effect caused or impeded by individual differences (for example innovativeness, expertise, position, experience, age, gender, experience with tools, computer self-efficacy, experience with topic of experiment).	Randomly assigned participants to treatment conditions. Measured individual differences during the pretest as a control.
Effect caused or impeded by experimenter expectations.	Evaluators did not know which tool was assigned to each individual. Video taped sessions to allow for assessment of differential treatment by the facilitator. Used standardized script for each experiment session.
Individuals are more prone to rate and rank their ideas higher than other group member's ideas because the group evaluation task is overly onerous.	Limited number of ideas generated and limited evaluation to rating only. Randomly sorted ideas for evaluation.
High ranking or aggressive individuals in the group experiment will influence how others rate and rank ideas.	Use of GSS anonymity features helped alleviate this issue. I also measured choice shift to determine whether this indeed occurred.
Individuals may come up with more radical and novel ideas if provided an incubation period. Thus, providing an incubation period after the initial search results would eliminate or reduce the effect.	In this case, Athens appears to reduce the need for an incubation period.
Individuals may require domain relevant knowledge in order to effectively find novel knowledge using Athens. Athens may not support individual learning otherwise.	This was not an issue for the field experiment, since the search task was based on an existing organizational problem/question. In the lab experiment, the topic chosen was hypothesized to be very familiar to students. Familiarity with the topic was measured in the pretest and posttest.

Table C.2 – Construct Validity

Type of Threat	Description	Mitigation?
Inadequate Explication of Constructs	Inadequately analyzing and defining the construct beforehand can result in a mismatch between operations and the intended construct.	Conducted card sorting exercise. Pilot tested experiment and questions.
Construct Confounding	Operations that really test two constructs, not the intended construct.	Conducted card sorting exercise. Pilot tested experiment and questions. Asked PhD students and faculty to review questions.
Mono-Operation Bias	Single operationalizations of a construct may underrepresent that construct.	Used multiple items to measure a construct or multiple measures.
Mono-Method Bias	When the method biases or influences the results.	Used multiple methods to measure the results. Lab and field experiment overlapped in terms of testing certain propositions. Also, asked for participant's perceptions and perceptions of five industry evaluators.
Confounding Constructs with Levels of Constructs	When only certain levels of the construct have been tested, rather than all levels of the intended construct.	Tested three distinct tools and three processes.
Reactive Self Report Changes	When prior to assignment participants answer pretest questions in such a way as to deliberately affect their assignment to conditions.	Pretest did not determine treatment assignment.
Reactivity to Experimental Situation	Sometimes called hypothesis guessing. When participants try to guess what the experimenter is doing and provide responses accordingly.	Asked student participants not to discuss the experiment with others. Asked the organizational sponsor not to reveal the study's hypotheses to participants. Left sensitive questions to posttest.
Experimenter Expectancies	When the researcher treats each treatment group differently, based on the results they expect.	Evaluators did not know which tool was assigned to each individual. Video taped sessions to allow for assessment of differential treatment by the facilitator. Used standardized script for each experiment session.
Novelty and Disruption Effects	When the treatment is innovative or disruptive, the results may be affected by the excitement and curiosity of participants about the novelty, or the anger towards the disruption.	The novelty and disruption effect was held constant in the field experiment, however may have affected results in the lab experiment for Athens vs. Google and Vivisimo.
Compensatory Equalization	When the experimenter tries to equalize the benefits across treatment groups, minimizing the differences and diluting the treatment.	Benefits from each treatment were very similar and not substantial.
Compensatory Rivalry	When participants in the less desirable control groups actually try harder and compete to look good because of their assignment to this group.	Participants were not aware that there were "better" treatment conditions. They were only familiar with their treatment.

Type of Threat	Description	Mitigation?
Resentful Demoralization	When participants in the less desirable control groups are resentful, which changes how they respond to the measures.	Participants were not aware that there were “better” treatment conditions. They were only familiar with their treatment.
Treatment Diffusion	When some respondents in the control groups receive some parts of the treatment, usually through interaction with individuals in the treatment group.	Asked student participants not to discuss the experiment with others.

Table C.3 – Statistical Conclusion Validity

Type of Threat	Description	Mitigation?
Low Statistical Power	If a study has low power, an effect may not be detected.	Sample size calculated to achieve power level of 80%.
Violated Assumptions of Statistical Tests	Violating the assumptions of statistical tests may result in inaccurate inferences about the relationship between the independent and dependent variables.	Reviewed assumptions for MANOVA and factor analysis.
Unreliability of Measures	Unreliability of measures results in increased error variance and inaccurate inferences about the relationship between the independent and dependent variables.	Assessed reliability of measures.
Restriction of Range	If there is insufficient range (variance) between independent variables in each treatment condition, an effect may not be detected.	Tested three distinct tools and three processes.
Unreliability of Treatment Implementation	If the treatment condition is not applied consistently to each participant, the effect of the treatment may be diminished.	Standard application, training, and setting. Used standardized script for each experiment session.
Extraneous Variance in Experiment Setting	Extraneous variance, or noise, related to the experimental setting increases the error variance.	Used Queen’s Executive Decision Centre (QEDC) lab, which helped to minimize extraneous variance (noise, disruptions, etc).
Heterogeneity of Units	The heterogeneity of participants within each group will make the effect harder to detect.	Students were used in the lab experiment – a fairly homogeneous group. However, some heterogeneity is desired so that a common trait doesn’t affect results. Participants from the same organization were used in the field experiment. Again, the common characteristic of “employment at the organization” means participants were not highly heterogeneous. Measured individual differences in control measurements.

Table C.4 – Internal Validity

Type of Threat	Description	Mitigation?
Ambiguous Temporal Precedence	If it is not clear in which order the variables occurred, causality may be difficult to establish	This threat is not applicable to the lab and field experiments.
Selection	When differences exist between the group of participants assigned to one treatment condition and those in another treatment condition, they could potentially explain differences on the outcomes between groups	Used random assignment.
History	Events, other than the treatment itself, that occur between the initiation of the treatment and posttest that could explain the outcome	Not enough time passed to make this a threat.
Maturation	Natural changes that occur between the initiation of the treatment and posttest, which could explain the outcome	Not enough time passed to make this a threat.
Regression to the Mean	When participants are selected based on scores that are lower or higher than average scores, they will tend to score closer to the average overtime. This trend is an alternative explanation for the outcome.	People not chosen based on scores.
Attrition	When participants that drop out of the study are systematically different from those that remain, the outcome results may be the result of these differences, not the treatment.	Used control questions to measure differences between participants that leave and those that remain.
Testing	When the participants are tested using the same test multiple times, their scores may improve as the result of familiarity and practice with the test, not the result of the treatment.	In lab experiment – subjects tested only once. In field experiment – subjects tested multiple times. But this was part of the treatment condition.
Instrumentation	When changes occur in the measurement instrument over time, the outcome could be the result of the instrument changes, not the treatment.	It is possible that the evaluators who evaluated the ideas from the lab experiment could become tired. This could affect measurement. Measured inter-rater reliability. Limited fatigue by breaking up ideas into groups and asked evaluators to take a break between groups of ideas. Provided each judge with the ideas sorted in a different order.

APPENDIX D CARD SORT

Table D.1 – Round 1: Mapping of Categories Identified by Judges to Intended Constructs

Intended Construct	Matching Category Labels identified by Judges¹
Overall radicalness	How does the idea relate to what is currently offered by the company; Innovativeness of idea; Novelty of idea
New output	How will the company react to this idea; Idea's influence on new offerings by firm; Changes resulting from the new idea; New product from idea
New processing	How will the idea impact existing technologies and systems of the company; Newness introduced by idea AND The idea's fit with the existing organization (TANGIBLES); Potential resource disruption from the idea AND Disruption
Relevant radical/interesting	How will the company react to this idea; The value of the idea to the organization
Overall Benefit	How will the idea impact the overall performance of the company; Idea's impact on firm competitiveness; The value of the idea to the organization; Competitive advantage or benefits
Cost/benefits analysis	How will the idea impact the overall performance of the company; Idea's impact on firm competitiveness; The value of the idea to the organization; Competitive advantage or benefits
Ethical analysis	Idea's degree of fit or compatibility with organization; Idea's degree of disruptiveness; The idea's fit with the existing organization (INTANGIBLES); Image -reactions from organizational outsiders
Protection of organizational strengths	Idea's degree of fit or compatibility with organization; The idea's fit with the existing organization (INTANGIBLES)
Compatibility	How will the idea impact the overall company structure and operations; Idea's degree of fit or compatibility with organization AND Idea's degree of disruptiveness; The idea's FIT with the existing organization (TANGIBLES) AND The idea's FIT with the existing organization (INTANGIBLES); Potential resource disruption from the idea AND Disruption caused by idea on organizational culture/processes/structure (intangible part of organization)
Stakeholder analysis	How would the company's stakeholders respond to the idea; Reaction of firm's main stakeholders to idea; Reactions to the idea; Image -reactions from organizational outsiders
Future position	How will the idea impact the overall performance of the company; Idea's impact on firm competitiveness; Why idea is important for the future; Competitive necessity
Deeper knowledge	In what way did the tool impact how much I know about the subject; Incremental knowledge acquired with tool; The Tool - gave me an amount (some) knowledge on top of what I already have; Increase on depth of knowledge
Broader knowledge	In what way did the tool impact how much I know about the subject; Incremental knowledge acquired with tool; The Tool - gave me an amount (some) knowledge on top of what I already have; Learning of relationships with similar concepts AND Increase on knowledge scope, breadth of knowledge
Novel knowledge	In what way did the tool impact how much I know about the subject; Unexpected new knowledge acquired with tool AND New knowledge acquired with tool; The Tool - "radicalized" the way I look at a subject; Discovery of previously unknown relation between topic and other results
Information foraging	How did the tool function; Tool's degree of accuracy; Adequacy of results
Interpretation	How did the tool function; Level of refinement of results; Meaningful results

¹ “AND” identifies multiple labels provided by the same judge. Items from the intended construct were placed in these multiple labeled categories.

Intended Construct	Matching Category Labels identified by Judges¹
	(not just the data)
Intuit	How did the tool function; Tool as a source of novel insights; The Tool - "radicalized" the way I look at a subject AND How meaningful it is. What it does above and beyond simply getting results; Development of new perspectives about the topic
Search term definition	How did the tool function; Level of effort required on part of user; How to use the tool; Implementation of search
Mental model tuning	In what way did the tool impact my understanding of the subject AND In what way did the tool impact how I think about the subject; Tool as a source of knowledge validation; The Tool - confirmed knowledge. In the same place where I started. Verification; Confirmation of information already known
Mental model maintenance	In what way did the tool impact my understanding of the subject AND In what way did the tool impact how I think about the subject; Incremental knowledge acquired with tool; The Tool - gave me an amount (some) knowledge on top of what I already have; Increase in understanding (in general)
Mental model building	In what way did the tool impact my understanding of the subject AND In what way did the tool impact how I think about the subject; Tool as a source of novel insights; The Tool - "radicalized" the way I look at a subject; Development of new perspectives about the topic

Table D.2 – Round 1: Reworded Items

Item	Original wording	New wording	Issue
RAD1	This idea is radical.	This idea is radically new to the organization.	Doesn't fit
INT1	This idea would be of interest to the organization.	This idea will interest the organization.	Doesn't fit and confusing
INT1	This idea may be considered for follow-up by the organization.	This idea will likely be pursued by the organization.	Doesn't fit and confusing
NOVEL1	This tool helped me learn concepts of which I was not previously aware.	This tool helped me learn concepts that are new to me.	Loaded onto the "Broad knowledge" category.
INTUIT1	The tool pointed me towards results that were interesting to me.	The tool helped me to develop interesting insights from the results.	Loaded onto the "Information Foraging" category.
MMM1	The tool helped me update my knowledge of this subject.	The tool helped me expand my knowledge of this subject.	Loaded onto "Tuning" and "Deep knowledge"
MMB5	The tool expanded the scope of how I think about this subject.	The tool changed the scope of how I think about this subject.	Loaded onto "Deep knowledge"
MMB6	The tool made me question my preconceptions about this subject	The tool made me reconsider my overall views of this subject.	Loaded onto "Tuning"

Table D.3 – Round 2: Mapping of Categories Identified by Judges to Intended Constructs

Intended Construct	Matching Category Labels identified by Judges
Overall radicalness	Novelty of idea
New output	Company growth and sustainability outcomes
New processing	Critical success factors for implementation
Relevant radical/interesting	Attractiveness/appeal of idea
Overall benefit	Attractiveness/appeal of idea
Cost/benefits analysis	Company growth and sustainability outcomes
Ethical analysis	Ease of implementing idea
Protection of organizational strengths	Ease of implementing idea AND Attractiveness/appeal of idea
Compatibility	Ease of implementing idea
Stakeholder analysis	Impact of implementing idea (on stakeholders); Stakeholder reaction
Future position	Company growth and sustainability outcomes; Possible strategic impact of the idea
Deeper knowledge	Ability to increase information/learning about a familiar topic; Substantiated individuals existing knowledge
Broader knowledge	Ability to increase information/learning about a familiar topic; Created new knowledge
Novel knowledge	Ability to increase information/learning about a new topic; Created new knowledge
Information foraging	Ease of use for info searching; Information search (search process)
Interpretation	Tool's ability to process information; Information presentation and selection
Intuit	Tool's ability to facilitate lateral thinking and integration; Information presentation and selection
Search term definition	Ease of use for info searching; Information search (search process)
Mental model tuning	Ability to confirm current level of information/knowledge about a familiar topic; Substantiated individuals existing knowledge
Mental model maintenance	Ability to increase information/learning about a familiar topic; Substantiated individuals existing knowledge
Mental model building	Tool's ability to broaden perspective (open minds); Created new knowledge

Table D.4 – Round 2: Reworded Items

Item	Original wording	New wording	Issue
MMM4	The tool broadened my understanding of this subject.	The tool extended my understanding of this subject.	Loaded onto "Tuning" and "Deep knowledge"

Table D.5 – Round 3: Item Issues

Item	Original wording	Issues	Resolution
STKHD5	This idea will not negatively affect the organization's reputation.	Originally part of ethical analysis but grouped with stakeholder reactions in rounds 1 and 2. Loaded onto compatibility in Round 3.	Remove item.
CMPINTAN3	This idea will positively affect the organization's existing strengths.	Originally part of protection of organizational strengths, but grouped with compatibility in rounds 1 and 2. Loaded onto overall benefits in round 3.	Take away distinction between intangibles and tangibles. This distinction is confusing for participants.
STKHD6	This idea will interest the organization.	Part of radical relevant/interesting, but grouped with stakeholder reaction after rounds 1 and 2. This item was still having loading problems in round 3. Loaded to future, radical and doesn't fit.	Create new category: Level of organization's interest in the idea.
STKHD7	This idea will likely be pursued by the organization.	Part of radical relevant/interesting, but grouped with stakeholder reaction after rounds 1 and 2. This item was still having loading problems in round 3. Loaded to future, radical and doesn't fit.	Create new category: Level of organization's interest in the idea.
NOVEL4	This tool helped me learn about associations that surprised me.	Consistently not grouped into novel knowledge category.	Reword: This tool helped me learn novel associations to this topic that surprised me.

Table D.6 – Round 4: Item Issues

Item	Original wording	Issues	Resolution
BROAD1	This tool helped me learn additional aspects about the topic.	Only loaded once to BROAD category. However, most cross loadings were to Deep, a highly similar category.	Left as is. Reexamined after round 5.
BROAD2	This tool helped me learn additional concepts related to the topic.	Only loaded twice to BROAD category.	Left as is. Reexamined after round 5.
BROAD4	This tool helped me learn relationships to other similar concepts.	Only loaded twice to BROAD category.	Left as is. Reexamined after round 5.
DEEP1	This tool helped me learn more things about what I already know.	Loaded three times to DEEP but also loaded three times to TUN.	Noted as an item to review after round 5.
DEEP2	This tool helped me learn incremental details about what I already know.	Loaded three times to DEEP but also loaded three times to TUN.	Noted as an item to review after round 5.
NOVEL2	This tool helped me learn concepts that I didn't know were related to the search terms provided.	Loaded three times to NOVEL, but had inconsistent cross loadings to other categories.	Reworded to make this item less confusing: The tool helped me learn new concepts that I had thought were unrelated to the subject.
TUN1	The tool helped me expand my knowledge of this subject.	Consistently loaded onto BROAD.	Moved to the BROAD category as BROAD5.
TUN3	The tool increased my understanding of this subject.	Did not load to TUN at all. The cross loadings were to other categories.	Item wording reviewed, however it did reflect the category. Thus, left as is. Reexamined after round 5.
TUN4	The tool extended my understanding of this subject.	Only loaded onto TUN once.	Left as is. Reexamined after round 5.
TUN5	The tool helped me refine my knowledge of this subject.	Only loaded onto TUN once.	Reworded to: The tool helped me refine my understanding of this subject.
INTUIT4	The tool helped me to develop new insights about this subject.	Only loaded onto INTUIT once with many cross loadings across categories.	Reworded to: The tool helped me to develop insights about this subject.

Table D.7 – Round 5: Item Issues

Item	Issues	Resolution
BROAD1	Only loaded once to BROAD category. However, most cross loadings were to Deep, a highly similar category.	Removed item after round 5.
BROAD2	Only loaded twice to BROAD category.	Removed to reduce number of items after round 5.
BROAD4	Only loaded twice to BROAD category.	Retained.
DEEP1	Loaded three times to DEEP but also loaded three times to TUN.	Removed after round 5 since it only loaded to DEEP once in round 5.
DEEP2	Loaded three times to DEEP but also loaded three times to TUN.	Retained as is because it loaded twice to DEEP in round 5.
NOVEL2	Loaded three times to NOVEL, but had inconsistent cross loadings to other categories.	Retained.
TUN1	Consistently loaded onto BROAD.	Retained as BROAD5.
TUN3	Did not load to TUN at all. The cross loadings were to other categories.	Retained.
TUN4	Only loaded onto TUN once.	Retained.
TUN5	Only loaded onto TUN once.	Retained.
INTUIT4	Only loaded onto INTUIT once with many cross loadings across categories.	Retained.

Table D.8 – Summary of Items Dropped and Retained

Construct	Items Dropped for Posttest Survey	Items Retained for Posttest Survey ²
Ben	Dropped BEN3. Used in idea evaluation survey.	Retained BEN1, BEN2.
Cmpt	Dropped CMPT1 through 8.	
Future	Dropped FUTURE 1 and 2. FUTURE2 Used in idea evaluation survey.	
Int	Dropped INT1.	Retained INT2. Also used in idea evaluation survey.
Stkhd	Dropped STKHD2 through 4.	Retained STKHD1
NewO	None.	Retained NEWO1, NEWO2. Combination used in idea evaluation survey.
NewP	Dropped all items.	
Rad	Dropped RAD3. Used in idea evaluation survey.	Retained RAD1, RAD2. Combination used in idea evaluation survey.
Deep	Dropped DEEP1. Note, DEEP1 was identified as a possible problem in round 4.	Retained DEEP2, DEEP3, DEEP4.
Broad	Dropped BROAD1 and BROAD2. Note, BROAD1 was identified as a possible problem in round 4.	Retained BROAD3, BROAD4, BROAD 5.
Novel	None.	Retained NOVEL1, NOVEL2, NOVEL3, NOVEL4. Retained all 4 items since this is a critical construct being measured.
MMM	Dropped MMM2 and MMM6. All items loaded well (close to 100%), but MMM2 and MMM6 appeared to duplicate some of the other items. Thus, these items were removed in order to limit the size of the survey.	Retained MMM1, MMM3, MMM4, MMM5.
TUN	None.	Retained TUN2, TUN3, TUN4, TUN5. Kept all 4 items to mirror the items kept for the other learning mode constructs.
MMB	Dropped MMB1, and MMB5.	Retained MMB3, MMB4, MMB6, MMB7. Kept the best loading items.
Forage	None.	Retained FORAGE1, FORAGE2. Both items retained since there are only two items
Sterm	None.	Retained STERM1, STERM2. Both items retained since there are only two items.
Intuit	Dropped INTUIT2. INTUIT1, INTUIT3, and INTUIT4 all seemed to tap into insight development whereas INTUIT2 referred to patterns and, thus, possibly a separate construct.	Retained INTUIT1, INTUIT 3, INTUIT 4.
Interp	None.	Retained INTERP1, INTERP2, INTERP3. All three items retained since there are only three items.

² And idea evaluation survey where applicable.

Table D.9 – Benefit Items Dropped and Retained

Item	Item Wording	Category Label	Status
BEN1	This idea will be potentially beneficial.	Overall benefits	Retained
BEN2	This idea will be very valuable to the organization.	Overall benefits	Retained
BEN3	This idea will result in financial benefits that outweigh the financial costs.	Overall benefits	Used in idea evaluation survey only
CMPT1	This idea is congruent with the organization's culture.	Compatibility with organization	Discarded
CMPT2	This idea will be compatible with the organization's capabilities.	Compatibility with organization	Discarded
CMPT3	This idea will positively affect the organization's existing strengths.	Compatibility with organization	Discarded
CMPT4	This idea will not harm the organization's existing strengths.	Compatibility with organization	Discarded
CMPT5	This idea is consistent with the organization's strategy.	Compatibility with organization	Discarded
CMPT6	This idea will not disrupt the organization's structure.	Compatibility with organization	Discarded
CMPT7	This idea will not disrupt the organization's systems.	Compatibility with organization	Discarded
CMPT8	This idea will be technically feasible to implement based on the organization's current technological competencies and knowledge-base.	Compatibility with organization	Discarded
FUTURE1	This idea is required for the long-term survival of the organization.	Importance for future outlook	Discarded
FUTURE2	This idea is required for the competitive success of the organization.	Importance for future outlook	Used in idea evaluation survey only
INT1	This idea will interest the organization.	Level of organization's interest in the idea	Discarded
INT2	This idea will likely be pursued by the organization.	Level of organization's interest in the idea	Retained and used in idea evaluation survey
STKHD1	This idea will result in a positive reaction from the organization's customers.	Stakeholder reaction	Retained
STKHD2	This idea will result in a positive reaction from the organization's investors.	Stakeholder reaction	Discarded
STKHD3	This idea will result in a positive reaction from the organization's suppliers.	Stakeholder reaction	Discarded

Item	Item Wording	Category Label	Status
STKHD4	This idea will result in a positive reaction from the organization's employees.	Stakeholder reaction	Discarded

Table D.10 – Radicalness Items Dropped and Retained

Item	Item Wording	Category Label	Status
NEWO1	This idea would result in the development of a new product.	New outputs	Retained and used in idea evaluation survey
NEWO2	This idea would result in offering a new service.	New outputs	Retained and used in idea evaluation survey
NEWP1	This idea would require using new processes at the organization.	Change to firm's internal production processing	Discarded
NEWP2	This idea would require using new materials at the organization.	Change to firm's internal production processing	Discarded
NEWP3	This idea would require using new equipment at the organization.	Change to firm's internal production processing	Discarded
RAD1	This idea is radically new to the organization.	Overall novelty of the idea	Retained and used in idea evaluation survey
RAD2	This idea is novel to the organization.	Overall novelty of the idea	Retained and used in idea evaluation survey
RAD3	This idea is substantially different from existing technologies.	Overall novelty of the idea	Retained and used in idea evaluation survey

Table D.11 – Knowledge Discovery Items Dropped and Retained

Item	Item Wording	Category Label	Status
DEEP1	This tool helped me learn more things about what I already know.	Acquired deeper knowledge	Discarded
DEEP2	This tool helped me learn incremental details about what I already know.	Acquired deeper knowledge	Retained
DEEP3	This tool helped me learn deeper knowledge about the subject.	Acquired deeper knowledge	Retained
DEEP4	This tool helped me learn more specific details about what I already know.	Acquired deeper knowledge	Retained
BROAD1	This tool helped me learn additional aspects about the topic.	Acquired broader knowledge	Discarded
BROAD2	This tool helped me learn additional concepts related to the topic.	Acquired broader knowledge	Discarded
BROAD3	This tool helped me develop a broader understanding of a topic with which I am already familiar.	Acquired broader knowledge	Retained
BROAD4	This tool helped me learn relationships to other similar concepts.	Acquired broader knowledge	Retained
BROAD5	The tool helped me expand my knowledge of this subject.	Acquired broader knowledge	Retained
NOVEL1	This tool helped me learn concepts that are new to me.	Acquired novel knowledge	Retained
NOVEL2	The tool helped me learn new concepts that I had thought were unrelated to the subject.	Acquired novel knowledge	Retained
NOVEL3	This tool helped me learn about new knowledge previously unknown to me	Acquired novel knowledge	Retained
NOVEL4	This tool helped me learn novel associations to the topic that surprised me.	Acquired novel knowledge	Retained

Table D.12 – Learning Mode Items Dropped and Retained

Item	Item Wording	Category Label	Status
MMM1	The tool helped me confirm my understanding of the subject.	Confirm current knowledge/mental models	Retained
MMM2	The tool helped me verify my knowledge of the subject.	Confirm current knowledge/mental models	Discarded
MMM3	The tool helped me maintain my perspective regarding this subject	Confirm current knowledge/mental models	Retained
MMM4	The tool supported my current understanding of this subject	Confirm current knowledge/mental models	Retained
MMM5	The tool reinforced my understanding of this subject.	Confirm current knowledge/mental models	Retained
MMM6	The tool validated my understanding of this subject.	Confirm current knowledge/mental models	Discarded
TUN2	The tool helped me augment my current understanding of this subject.	Update current knowledge/mental models	Retained
TUN3	The tool increased my understanding of this subject.	Update current knowledge/mental models	Retained
TUN4	The tool extended my understanding of this subject.	Update current knowledge/mental models	Retained
TUN5	The tool helped me refine my understanding of this subject.	Update current knowledge/mental models	Retained
MMB1	The tool challenged me to understand the subject in a new light.	Reconfiguring current knowledge/mental models	Discarded
MMB3	The tool changed my perspective on this subject.	Reconfiguring current knowledge/mental models	Retained
MMB4	The tool re-oriented my thinking on this subject	Reconfiguring current knowledge/mental models	Retained
MMB5	The tool changed the scope of how I think about this subject.	Reconfiguring current knowledge/mental models	Discarded
MMB6	The tool made me reconsider my overall views of this subject.	Reconfiguring current knowledge/mental models	Retained
MMB7	The tool helped me reframe my thinking regarding this concept.	Reconfiguring current knowledge/mental models	Retained

Table D.13 – Learning Process Items Dropped and Retained

Item	Item Wording	Category Label	Status
STERM1	With this tool, I had to refine my search terms several times in order to get the results I wanted.	Search term refinement	Retained
STERM2	The tool required me to provide precise terminology in the search terms.	Search term refinement	Retained
FORAGE1	The tool found the information for which I was looking.	Learning support - Search capabilities	Retained
FORAGE2	The tool provided me with access to the information for which I was searching.	Learning support - Search capabilities	Retained
INTUIT1	The tool helped me to develop interesting insights from the results.	Learning support - developing insights from patterns in results	Retained
INTUIT2	The tool helped me to discover patterns in the results.	Learning support - developing insights from patterns in results	Discarded
INTUIT3	Because of the interesting results it provided, the tool helped me to develop my own insights that I had not thought about previously.	Learning support - developing insights from patterns in results	Retained
INTUIT4	The tool helped me to develop insights about this subject	Learning support - developing insights from patterns in results	Retained
INTERP1	The tool interpreted the raw results for me.	Learning support - developing meaning in the results	Retained
INTERP2	The tool provided me with refined and meaningful results.	Learning support - developing meaning in the results	Retained
INTERP3	The tool provided me with a summary of the results rather than all of the raw results.	Learning support - developing meaning in the results	Retained

APPENDIX E
QUESTIONNAIRES USED IN THE LAB EXPERIMENT

Novel Knowledge Experiment - Part 1

E.1 Pretest Questionnaire

Letter of Information and Informed Consent

Thank you for considering participation in my research study "Using IT to Support the Discovery of Novel Knowledge in Organizations". This study is part of my dissertation as a PhD student at Queen's University, School of Business. This letter outlines the details of your participation.

The topic being researched is the discovery of novel knowledge on the Web. Novel knowledge is defined as knowledge that is not currently known to the organization, interesting, relevant, but indirectly connected to what the organization currently knows and therefore difficult to find.

The main purpose of this study is to evaluate the effectiveness and usability of different tools for discovering novel knowledge on the Web. As a part of this research, you will be asked to complete this brief questionnaire (part 1) and participate in an experiment in the Queen's Executive Decision Centre (part 2), where you will be provided a search task and asked to complete that task using a search tool pre-assigned to you. The following questionnaire will take approximately 15 minutes to complete; the experiment, including training and final questionnaire completion, will last approximately 75 minutes. Therefore, your total time commitment is 90 minutes. A few days after the experiment, you may be contacted by e-mail to answer specific questions for clarification purposes (for example, clarifying your feedback on the tools used).

A prize of \$100 will be awarded to the participant who most effectively completes the search task. Details about the effectiveness criteria will be explained at the start of the experiment. The winning participant will be notified via email 3 to 4 weeks following the completion of the experiment.

All sessions will be recorded on videotape. You will be asked for your permission to videotape your sessions on a consent form provided to you at the start of the experiment. These tapes will only be accessible to the researcher, Tracy Jenkin, and her supervisors, David Skillicorn and Yolande Chan.

The confidentiality of the information you provide will be protected by concealing your name. Your name, email address, student number and participant code will be kept in a separate document, to be used to contact you for clarification questions later and provide you with extra credit for your participation. Only your participant code (e.g. participant A) and assigned tool will be connected with the information you provide during the experiment. These documents and the videotape from the experiment will be kept safe in a locked file drawer in the Queen's University School of Business faculty offices. You will not be quoted in written form unless you give

Novel Knowledge Experiment - Part 1

explicit consent to be quoted on the consent form provided to you at the start of the experiment.

There are no known physical, psychological, economic, or social risks associated with this study. Your participation in this study is completely voluntary and you may withdraw from this study at any time, for any reason with no effect on your standing at Queen's University. You may also decline to answer any specific questions. The results of this research will be part of a PhD thesis that will be submitted to Queen's University. The academic community and any other person interested in these research results will have access to it through Queen's University. This research may also be published in the form of a journal article or conference proceeding at a later stage and thus, available to the general public or as a secondary source for other researchers. Your confidentiality is guaranteed. Your name will not be connected to any publication of these research results.

If you would like further information about the study, or have additional questions or concerns, please feel free to contact any one of the following:

Tracy Jenkin, PhD candidate tjenkin@business.queensu.ca

Dr. David Skillicorn skill@cs.queensu.ca

Dr. Yolande Chan ychan@business.queensu.ca

Dr. Bill Cooper Unit Research Ethics Board (613) 533-2333
bcooper@business.queensu.ca

Dr. Joan Stevenson General Research Ethics Board (613) 533-6288
stevensj@post.queensu.ca

Thank you for your time and interest.

Tracy Jenkin, PhD Candidate

Novel Knowledge Experiment - Part 1

By choosing to participate in this study, you are agreeing to the following statement: I consent to participate in the study "Using IT to Support the Discovery of Novel Knowledge in Organizations". I have read the Letter of Information, had any questions answered to my satisfaction and will keep a copy of this letter for my records. I confirm that I understand the provisions around confidentiality to protect my identity. I also understand that my participation is voluntary and I have been told who to contact if I have questions/concerns about this study.

1. Do you consent to participate in this study?

 I consent to participate in this study
 I do NOT consent to participate in this study

2. Please enter your student number:
3. Please enter your first name:
4. Please enter your last name:
5. Please enter your QLINK email address:
6. Gender: Male Female
7. Current age:
8. Is English your first language? Yes No
9. Program enrolled in at Queen's:
 Commerce
 Arts and Science
 Engineering
 Other (Please specify):

10. Current year in your program at Queen's:
 First year
 Second year
 Third year
 Fourth year
 Other (Please specify):

Novel Knowledge Experiment - Part 1

11. In general:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I am familiar with the topic of MP3 technology.						
0	0	0	0	0	0	0
I am familiar with the Google search tool.						
0	0	0	0	0	0	0
I use Google frequently.						
0	0	0	0	0	0	0
I am familiar with the Vivisimo search tool.						
0	0	0	0	0	0	0
I use Vivisimo frequently.						
0	0	0	0	0	0	0

12. In general:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I feel confident searching the Web for new knowledge.						
0	0	0	0	0	0	0
I feel confident finding information on the Web using a search engine.						
0	0	0	0	0	0	0
I feel confident reviewing the search engine results page						
0	0	0	0	0	0	0
I feel confident navigating the search engine results page by following hyperlinks.						
0	0	0	0	0	0	0

This part of the questionnaire asks you about your ability to use an unfamiliar piece of software. Often in our jobs we are told about software packages that are available to make work easier. For the following questions, imagine that you were given a new software package for some aspect of your work. It doesn't matter specifically what this software package does, only that it is intended to make your job easier and that you have never used it before.

Novel Knowledge Experiment - Part 1

The following questions ask you to indicate whether you could use this unfamiliar software package under a variety of conditions. For each of the conditions, please indicate your level of confidence in your ability to complete the job using the software package. Rate your confidence level by selecting a number from 1 to 10, where 1 indicates "Not at all confident", 5 indicates "Moderately confident", and 10 indicates "Totally confident".

13. Under the following conditions, I could complete the job using the software package...

Level of confidence:

(1) Not at all confident	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10) Totally confident
If there was no one around to tell me what to do as I go.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
If I had never used a package like it before.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
If I could call someone for help if I got stuck.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
If someone else had helped me get started.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
If I had just the built-in help facility for assistance.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
If I had used similar packages before this one to do the same job.									
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

During the experiment, you will be asked to complete the search task described below. In order to complete this task you will need to choose descriptive terms, which will be provided to a search tool. These terms should be generic, standard terms for the topic area you are searching. In practice, two terms is ideal. Each 'term' can be composed of one, two, or three words, creating a phrase. For example, "digital imagery" is a single term. Each term will be provided to the search tool in "quotes". For example, if digital imagery was one term and electronics the second term, the query would be "digital imagery" + "electronics".

Keep the scope of your terms to a reasonable size. Highly specialized terms may not provide interesting results. Terms that are too general may also be problematic and hide many interesting results. For example, if one was conducting a search in a medical context, the term "cancer cure" is too general; the term "uranium hexachloride ingestion" is too specific. Terms such as "Raynaud's disease" and "blood", or "carpal tunnel syndrome" and "repetitive strain injury" work much better.

Novel Knowledge Experiment - Part 1

Additional instructions will be provided when the experiment starts. Please read the task and answer the questions that follow carefully.

Your Task:

You work for a company that designs, develops and sells MP3 players. Thus, your company knows and understands MP3 technology and the market for MP3 players very well. They would like to find **novel** ways of applying their MP3 expertise – for example, to develop new services, new products, or new markets. During the experiment, you will be asked to use a specific search tool to search the Web for new potential applications of your company's MP3 knowledge, record these ideas, and then evaluate them on how **novel** and how **beneficial** they are to your company.

Below, please identify two sets of descriptive terms that you would like to use in the search tool. You will use these terms during the experimental task, so please answer thoughtfully.

As a **practical example**, in the 90's Sony applied its extensive knowledge of electronics and digital devices to create a new product and market – digital cameras. In Sony's case, the descriptive terms could have been: 1) "digital imagery" and 2) "electronics". Sony's expertise in these areas lead them to apply their knowledge to digital cameras and photography, a novel product and market.

14. Descriptive Term 1:

15. Descriptive Term 2:

Thank you for completing **part 1** of this study. The next step is to complete the **experiment** in the Queen's Executive Decision Centre (**Goodes Hall, room 102 in the BASEMENT**) on the date and time you signed up for.



Ideas

E.2 Ideas Web Form

Please use this web form to document the ideas you generate during your search.

Document 4 to 5 ideas that you generate as you proceed with your search. Document the following: 1) the idea, 2) the webpage from which you generated or found the idea, and 3) the search terms you were using at the time you generated the idea.

Once you have documented an idea, click on the "next idea" button to document additional ideas.

1. Please enter your student number:
2. At which laptop station are you sitting?

Choose One...	V
---------------	---

3. Idea 1

Description:

4. Idea 1

Web Page link:

5. Idea 1

Search terms used:

Next Idea

6. Idea 2

Description:

7. Idea 2

Web Page link:

8. Idea 2

Search terms used:

Next Idea

Ideas

9. Idea 3

Description:

10. Idea 3

Web Page link:

11. Idea 3

Search terms used:

[Next Idea](#)

12. Idea 4

Description:

13. Idea 4

Web Page link:

14. Idea 4

Search terms used:

[Next Idea](#)

15. Idea 5

Description:

16. Idea 5

Web Page link:

17. Idea 5

Search terms used:

[Done](#)

Ideas

Thank you. Please raise your hand to let me know that you have completed this task.



Novel Knowledge Experiment - Survey 2

E.3 Posttest Questionnaire

In the following questionnaire, you will be asked to answer a number of questions about the ideas you generated today, your experience with the tool and the search task. In addition, you will be asked a few more questions about yourself.

Thank you for your time and interest.

Tracy Jenkin, PhD Candidate

Now that you have rank ordered your ideas, please take the top ranked idea from your list – in other words, the best idea that addresses the experimental task – and answer the following questions. Your answers will be provided to three independent judges who will evaluate your idea based on the novelty of the idea and potential benefit to the organization. The confidentiality of the idea you submit will be protected by concealing your name. As indicated in the information letter, a prize will be awarded to the student with the best idea, judged on the criteria just described.

1. Please briefly describe your top ranked idea:
2. Please explain the reasons why the organization should pursue this idea:
3. Did you have difficulty choosing which idea was the best?
 Yes
 No
4. If you answered **'yes'**, please provide comments on why you found it difficult to decide between ideas.
5. Please enter your student number:

Novel Knowledge Experiment - Survey 2

Please answer the following series of questions regarding your experience with the tool and the search task in today's experiment.

6. What type of knowledge did you discover with your tool?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
This tool helped me develop a broader understanding of a topic with which I am already familiar.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn incremental details about what I already know.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn concepts that are new to me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn novel associations to the topic that surprised me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn deeper knowledge about the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn relationships to other similar concepts.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me learn new concepts that I had thought were unrelated to the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me expand my knowledge of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn more specific details about what I already know.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn about new knowledge previously unknown to me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

7. How did this tool help you learn?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
The tool found the information for which I was looking.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with a summary of the results rather than all of the raw results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me to develop interesting insights from the results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with refined and meaningful results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
With this tool, I had to refine my search terms several times in order to get the results I wanted.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Because of the interesting results it provided, the tool helped me to develop my own insights that I had not thought about previously.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool interpreted the raw results for me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with access to the information for which I was searching.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me to develop insights about this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool required me to provide precise terminology in the search terms.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

8. How did this tool affect your understanding of this subject?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
The tool supported my current understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool changed my perspective on this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool increased my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool made me reconsider my overall views of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me maintain my perspective regarding this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me augment my current understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me reframe my thinking regarding this concept.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me confirm my understanding of the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool extended my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool reinforced my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me refine my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool re-oriented my thinking on this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

9. What were you trying to accomplish in your search?

10. By completing today's task...

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I increased my familiarity with the topic of MP3 technology.						
0	0	0	0	0	0	0

11. Describe the functionality of the tool you used for today's task:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
Searches the Web using keywords entered.						
0	0	0	0	0	0	0
Groups results into clusters of similar content.						
0	0	0	0	0	0	0
Draws attention to interesting and novel results that are indirectly related to keywords entered.						
0	0	0	0	0	0	0
Allows users to drill down into clusters of similar content.						
0	0	0	0	0	0	0
Includes links to results that are most similar to keywords entered.						
0	0	0	0	0	0	0
Results are indirectly connected to keywords, rather than most similar to keywords.						
0	0	0	0	0	0	0

12. Please answer the following questions regarding today's training:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
Sufficient training was provided at the beginning of the experiment.						
0	0	0	0	0	0	0
Additional training would have helped me complete the experimental task more effectively.						
0	0	0	0	0	0	0

Novel Knowledge Experiment - Survey 2

13. Please answer the following questions regarding the instructions provided today:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I understood what I was being asked to do in this experiment.						
0	0	0	0	0	0	0
The instructions provided by the facilitator were clear.						
0	0	0	0	0	0	0

14. Describe your satisfaction with using today's tool:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I found the tool useful for completing today's task.						
0	0	0	0	0	0	0
I found using the tool frustrating.						
0	0	0	0	0	0	0
I was better off using the tool to complete the task than without the tool.						
0	0	0	0	0	0	0

Please answer the following questions and provide examples from your search task today, where applicable:

15. What features of this tool assisted you with today's task?

16. What challenges did you encounter using this tool?

17. If this tool could be improved, what are the most important features that should be added or fixed?

Novel Knowledge Experiment - Survey 2

18. In general...

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
If I heard about a new information technology, I would look for ways to experiment with it.						
0	0	0	0	0	0	0
Among my peers, I am usually the first to try out new information technologies.						
0	0	0	0	0	0	0
In general, I am hesitant to try out new information technologies.						
0	0	0	0	0	0	0
I like to experiment with new information technologies.						
0	0	0	0	0	0	0

Thank you for completing survey 2 of this study. You have now completed all components of the experiment! Please do not discuss this experiment with your classmates in case they are also participating. Three to four weeks after the end of the entire experiment, I will email the winner of the \$100 gift certificate prize.

Thank you!
 Tracy Jenkin, PhD Candidate



Idea Evaluation

E.4 Ideas Evaluation Template

As discussed in the letter of information you were provided, in the following questionnaire you will be asked to evaluate several ideas generated during an experiment conducted this winter.

Thank you for your time.

Tracy Jenkin, PhD Candidate

By choosing to participate in this study, you are agreeing to the following statement: I consent to participate in the study "Using IT to Support the Discovery of Novel Knowledge in Organizations". I have read the Letter of Information, had any questions answered to my satisfaction and will keep a copy of this letter for my records. I confirm that I understand the provisions around confidentiality to protect my identity. I also understand that my participation is voluntary and I have been told who to contact if I have questions/concerns about this study.

1. Do you consent to participate in this study?
 I consent to participate in this study
 I do NOT consent to participate in this study

Idea Evaluation

Instructions for Idea Evaluation

As part of an experiment I conducted for my thesis research, individual participants were given the following task:

"You work for a company that designs, develops and sells MP3 players. Thus, your company knows and understands MP3 technology and the market for MP3 players very well. They would like to find **novel** ways of applying their MP3 expertise – for example, to develop new services, new products, or new markets.

As a practical example, in the 90's Sony applied its extensive knowledge of electronics and digital devices to create a new product and market – digital cameras.

Please use the search tool provided to search the Web for new potential applications of your company's MP3 knowledge."

Individuals were asked to generate ideas they believed were novel ways of applying their company's MP3 expertise, for example, to develop new services, new products, or new markets.

They were asked to document their best idea and told that this idea would be evaluated on how novel and how beneficial the idea is to the company.

The following questionnaire contains a subset of the ideas proposed by the participants to address the task above. Please review each idea – both the idea description and the idea rationale – and rate each idea on its novelty and benefit.

Idea # 1

Idea Description (what):

Idea Rationale and Importance (why):

Questions:

Note: N/A refers to "not applicable".

2. Using the following scale, how novel will this idea be to the organization, based on what the organization does now?:

(N/A) Not applicable	(1) Slightly new	(2)	(3)	(4) Moderately new	(5)	(6)	(7) Radically new
This idea will be...							
0	0	0	0	0	0	0	0

Idea Evaluation

3. Using the following scale, what will the impact of this idea be on the organization's products or services?:

(N/A) Not applicable	(1) Improve existing product functionality or services	(2)	(3)	(4) Add new functionality or dimensions to existing products/ services	(5)	(6)	(7) Create a new product or service
This idea will...							
0	0	0	0	0	0	0	0

4. Using the following scale, what will the impact of this idea be on the organization's technology?:

(N/A) Not applicable	(1) Use the company's existing technology	(2)	(3)	(4) Modify the company's existing technology	(5)	(6)	(7) Use a new technology in the company
This idea will...							
0	0	0	0	0	0	0	0

5. Using the following scale, how profitable will this idea likely be?:

(N/A) Not applicable	(1) Not profitable	(2)	(3)	(4) Marginally profitable	(5)	(6)	(7) Highly profitable
This idea will likely be...							
0	0	0	0	0	0	0	0

6. Using the following scale, how will this idea impact the organization's competitiveness?:

(N/A) Not applicable	(1) Not help the company achieve competitive success	(2)	(3)	(4) Help company achieve short-term competitive success	(5)	(6)	(7) Help company achieve long-term competitive success
This idea will likely...							
0	0	0	0	0	0	0	0

Idea Evaluation

7. Using the following scale, how likely is it that the organization will pursue this idea?:

(N/A) Not applicable	(1) Not at all likely	(2)	(3)	(4) Moderately likely	(5)	(6)	(7) Highly likely
The likelihood of the organization pursuing this idea is...							
0	0	0	0	0	0	0	0

Thank you for completing this questionnaire.

Tracy Jenkin, PhD Candidate



E.5 MP3 Summary Sheet

MP3 Technology and Industry Overview

MP3 Player Overview

An MP3 player, more generically called a digital audio player (DAP), is a device that stores, organizes and plays digital music files. MP3 is a common file compression format. However, there are other types of formats on the market; some that are proprietary, some that incorporate digital rights technology to restrict illegal copying, and some that are open.

There are two main types of MP3 players:

1. **Flash-based Players** - These are devices that hold MP3 files on internal or external media, such as memory cards. Due to technological limitations, these are relatively low-storage devices (up to 10GB). These devices are very resilient and do not suffer limitations that owners of hard drive-based players face, such as fears of dropping their player or fragmentation. Such players are commonly integrated into USB keydrives.
2. **Hard Drive-based Players or Digital Jukeboxes (HDD)** - Devices that read MP3 files from a hard drive. These players have higher capacities (1.5GB – 160GB) and thus thousands of songs can be stored in one MP3 player. Because of the storage capacity, devices that also display video and pictures are often hard drive-based.

In addition to ripped CD tracks, many MP3 players can accept downloaded music from online music stores. However, such stores often use proprietary formats and DRM (digital rights management), which means that the tracks can only be played on suitably licensed and restricted devices. For example, songs you purchase from iTunes might not work on any MP3 player other than an iPod.

History

MP3 players were first produced in 1997. The ease with which MP3 files could be copied created legal issues regarding copyrights of the artists and recording companies. The arrival of Apple Computer's iPod in 2001, combined with the opening of the iTunes Store in 2003 and thus creation of the legal music download business, greatly expanded the market. Since then, a large number of new MP3 players have been released, each promising to be an "iPod Killer".

In 2004, Microsoft introduced their Digital Rights Management (DRM) technology, which allows consumers to rent music from subscription music services such as Napster and Rhapsody and transfer it to their compatible MP3 players.

Market Trends

- Flash-based models continue to dominate the MP3 player product line due to increasing memory capacity and decreasing cost. According to the Topology

Research Institute, flash MP3 players account for about 66 percent of worldwide shipments and HDD-based units for 33 percent. This trend is expected to continue in the next few years.

- The design trend is toward larger screens and more defined display types.
- Flash MP3 players with 8GB up to 10GB storage capacity are being developed. Mainstream flash players come in 128MB to 2GB models.
- HDD-based players are available in 1.5GB up to 160GB versions.
- Smaller form factors and multi-function designs also continue to be popular. Units integrated with GPS and Wi-Fi should be available in 2007.
- Flash MP3 players have been incorporated into sunglasses, as demonstrated by the Oakley's "thump" model in 2004.
- The MP3 player industry remains strong, with forecast growth expected for several more years even as part of manufacturing and product development steadily moves toward the portable media player (PMP).
- Greater China manufacturers account for approximately 52 percent of global production.

Equipment

In general, MP3 players are portable and use headphones, although users often connect players to car and home stereos. Some MP3 players also include FM radio tuners and/or microphones for voice recording. Most players have semi-permanent rechargeable batteries while others have conventional battery bays for disposable or rechargeable batteries. In 2006, MSI developed and showcased the first solar powered player.

Apple's iPod devices incorporate a proprietary dock connector that allows them to connect to accessories such as chargers, speakers with built-in charging, or car players. Accessories that use the dock connector are only compatible with iPods.

The use of external flash cards makes it possible in principle to move collections of music between personal players, portable players with speakers, home audio systems, and car players, though in practice car and home players that handle MP3 files and folders properly are hard to find.

A number of manufacturers now produce Network MP3 players, which connect to a home Ethernet network, and receive an MP3 stream from a computer on the network. They are designed to connect to a home stereo, and are operated with a remote control.

A new type of MP3 player has emerged as a result of satellite radio companies' push into North American markets. Linked to a paid subscriber's account, these devices record a certain number of programs, allowing the user to listen to pre-recorded programs, popular songs or talk shows, while on the run.

New uses for MP3 players include podcasting, in which radio-like programs, or even TV-like video feeds, are automatically downloaded into the device to be played at the owner's convenience.

Current Issues

1. **Proprietary file formats and device compatibility.** Example - songs you purchase from iTunes might not work on any MP3 player other than an iPod.
2. **Copyrights, digital rights management (DRM) and the legality of circumventing DRM.** On the one hand, DRM protects the copyrights of artists by restricting illegal copying, but on the other hand, it infringes on the consumer's legitimate uses of the music files.
3. **Illegal cloning** of MP3 players in China.

References

This summary was compiled and adapted based on content found on the following sites:

- <http://www.chinasourcingreports.com/csr/Electronics/MP3-Players/p/CSRMPP/Executive-Summary.htm>
- http://en.wikipedia.org/wiki/Digital_audio_player

APPENDIX F

LAB EXPERIMENT – SCRIPT

F.1 Verbal Process Guidelines

Notes to Experimenter

As people come in, introduce myself. Get them seated at computers with the **lids down and papers on top of the computer (student name and list of search terms on top)**.

Slide 1 - Introduction

- Good morning/afternoon, my name is Tracy Jenkin. I am a PhD student here at Queen's School of Business and this experiment is part of my thesis research.
- In front of you, you will find a letter of information, which describes this experiment, and a letter of consent. This is the same information that you read in the online survey. The consent form does contain additional information such as consent for video taping and direct quotes.
- If you did not review the letter online, please do so now.
- This study is completely voluntary so if you wish to withdraw at any time, you are free to do so. And let me know if you do not consent to video taping. In that case I will not start the video camera.
- If you wish to participate, please sign the letter of consent and pass it to me before starting.
- And as indicated this experiment is scheduled to take 75 minutes.
- Once you have all completed the consent form, we will begin.

Slide 2 - Agenda

- Here is a brief overview of what we will be doing today.
- First I will introduce the task you will be completing. You also read about this task in your survey
- Next, I will provide some brief tool training on your assigned search tool: <insert tool name>.
- Next, you will begin your task.
- Once you have completed your task, I will give some brief training on Group Systems.
- Then you will answer some questions about your task, the search tool and your experience – some questions will be answered using the Group Systems Software. Most will be answered using an online survey.

Slide 3 - Task Intro

- Please review the task instructions (instructions at each computer terminal in paper form). When you have finished reading raise your hand.
- Over the next 30 to 40 minutes, you can use the tool - <insert tool name> - to complete the task you have been assigned. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.
- As per the instructions, please generate 4 to 5 ideas.
- As you proceed in your search, document the ideas you generate from your search in the Web form "Ideas" open on your desktop. Please document the idea, the website you were

looking at when you generated the idea, and the search terms you were using at the time when you generated the idea.

- If you are looking at the summary results page when you generate your idea, and it is the summary that prompts you to generate your idea, list that website address in the web page link.
- For the web link, just copy and paste the link from the Internet explorer into this form field.

For Athens Treatment Only

- Please document the idea, and the website you were looking at when you generated the idea. And in the search terms being used box, please enter the search terms listed at the top of the summary page. These were the search terms you provided to me in your online survey.

Slide 4 - Training

- Now I will give you some brief training on the <insert tool name>

For Google Treatment Only

- You may have used Google before, but for the benefit of anyone who has not used Google before, I am going to provide a quick overview.
- Slide 5: First you choose the search terms you want to try and enter them in the search keyword box. Then click search.
- Slide 6: You can use more advanced searching by pairing two terms together – in this case carbon fiber as a term and cello as a term together. Note the quotation marks and the plus sign.
- Slide 7: You can also use the advanced search feature.
- Slide 8: The results of the search will be displayed in a list, such as this, including links and metadata about the link.
- You can click on the links to see what each page is about.
- And you can conduct as many searches as you like – sometimes the terms you use don't give you what you want, so you can try new search terms.
- For the first search, use terms you provided in your survey

Now please begin completing your task. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.

For Vivisimo Treatment Only

- You may have used Vivisimo before, but for the benefit of anyone who has not used Vivisimo before, I am going to provide a quick overview.
- Slide 5: First you choose the search terms you want to try and enter them in the search keyword box. Then click search. Note – do not try Clusty. Only use Vivisimo.
- Slide 6: You can use more advanced searching by pairing two terms together – in this case carbon fiber as a term and cello as a term together. Note the quotation marks and the plus sign.
- Slide 7: You can also use the advanced search feature.
- Slide 8: The results of the search will be displayed in a list, such as this, including links and metadata about the link.
- Notice also that there are clustered results on the left, which are basically the categories of the returned results. So you can click on a category and see only the results in that

category, or just review and click the links in the body (of all results). To get back to the see all results, just click on the top category.

- You can click on the links to see what each page is about.
- And you can conduct as many searches as you like – sometimes the terms you use don't give you what you want, so you can try new search terms.
- For the first search, use terms you provided in your survey

Now please begin completing your task. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.

For Athens Treatment Only

- Athens is a prototype tool designed to find novel knowledge, developed here at Queen's.
- Athens searches the web, based on the keywords you provide, and finds knowledge that is related but indirectly connected to those search terms to help you discover novel knowledge.
- This is very different from what most search tools do. Most search tools report back to you highly related and connected terms. So as you are using the tool, keep this in mind. In other words, keep in mind that the tool is trying to help you find new related yet indirectly connected areas of knowledge.
- For example, I did a search on daily aspirin and benefits. Ordinarily daily aspirin is known for its impact on cardiac health. However, what I was finding was content related to cancer and diabetes. So perhaps there is a connection or a possible future connection.
- Because of the algorithms that Athens uses, it takes a number of hours to run a query. So I have pre-run each of your search terms that you provided me through Athens and loaded them on your desktop.
- You will search through these search results during your task.
- Slide 5: This is what the Athens Results Summary or Home page looks like.
- Athens groups or clusters the results into two hierarchical levels. The first level is the parent level. So Athens identifies 5 major clusters of results, in this case, that relate indirectly to the search term. The terms beside each cluster are the nouns that describe that cluster.
- If you click on the cluster name or scroll down, you will see that each major cluster contains a number of additional child clusters called novel clusters.
- You can access these novel clusters by clicking on the major cluster links.
- If you scroll down you can see a summary of the novel clusters and their relationship to the main clusters.
- Slide 6: So here are the novel clusters that make up major cluster 5. You can see there are nouns that describe this cluster.
- The web links associated with each novel cluster are the web pages Athens found and grouped together, based on similarity, to form this cluster. So they relate indirectly to the search term and are described jointly by these terms.
- So if you think this may be an interesting cluster to investigate to help you generate novel ideas, click on these links and see what the page is talking about.
- Note - It is helpful to click on a least a couple of the links in a cluster to get a better understanding of what the cluster is generally about.
- Feel free to jump around from novel cluster to novel cluster and between major parent clusters.
- Clicking on a link opens a new window or tab. So you can always find this Athens window again. In your web browser – Mozilla Firefox – a new tab is opened.

- If you find a broken link, just ignore and keep searching through the results.
- Slide 7: You will notice descriptors at the top of the novel clusters pages. These tags are nouns that are prominent across the clusters. The larger the tag, the more prominently this tag appears across these web pages.
- If you see a tag term that interests you, you can click on it here and it will take you to the place where it shows you which novel clusters are associated with it.
- Just a different way of thinking about the novel clusters. There is no right or wrong way to explore the Athens output.
- So as you can see, there are many ways of maneuvering and exploring the results.
- Slide 8: Some quick notes on navigation through Athens. From the Summary Athens Page, you can get to the novel clusters page by clicking on the novel cluster link. To get to one of the web pages in the novel cluster, click on the web link. Again, this opens a new tab in your browser. To get back to Athens, simply click on the “clustering information” tab at the top of the browser. And to get to the Athens Summary page from the Novel clusters page, simply use the browser’s back button.

Now please begin completing your task. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.

Notes to Experimenter

- At 30-minute mark, alert the group that 30 minutes has passed.
- Copy ideas into GSS from survey tool into rating. Then copy into ranking.

Slide 9 – Group Systems

- Please close your search tool and any other open document on your desktop
- To rank and rate the ideas you generated during your search task, we will use a software tool called Group Systems.
- Group Systems is primarily used for group meetings to help with brainstorming and collaboration.
- However today, we won’t be using the group capabilities of the tool, just the voting mechanisms for ranking and rating your ideas.
- To start the Group Systems software, you will dbl click on the colourful icon on the leftside of you desktop. Please do so now.
- Slide 10: You should see the following screen.
- Slide 11: When I start your session, you will be given some instructions and asked to rate your ideas according to how radical or novel they are and how potentially beneficial they are. You will see the question on top, and the idea on the bottom. You will indicate the degree to which you agree or disagree with each statement for that idea.
- Slide 12: Once you have answered a question, click next to proceed to the next question. Notice that as you move through the screens, you will answer a series of questions about one idea. So here we see a new question, but same idea. Once you have answered all the questions about an idea, you move onto the next idea and answer the same set of questions. So read the screen carefully so that you know which question and which idea is the subject.
- Slide 13: When you have completed the ratings, you will notice that the next button is no longer enabled. And you will be on the last question for the last idea.
- Click the ‘cast ballot button’, which is the ballot box icon up top.
- You will see the Group systems screen. At this point raise your hand so that I know you are ready to move to the next part.

- Slide 14: Next you will proceed to ranking. You will be given instructions and then allowed to rank your ideas by clicking and dragging your ideas to where you think they belong in the rank order. When you are moving an idea, you will notice the tow truck icon.
- Slide 15: when you have finished ranking, please cast your vote.

REMEMBER which idea you ranked as your top idea. You will need this information for the survey.

Final instructions after Rating and Ranking

Please complete survey – survey two on desktop. And please make sure you enter your student number. Once you have completed the survey, you are free to leave. Thank you for participating.

APPENDIX G

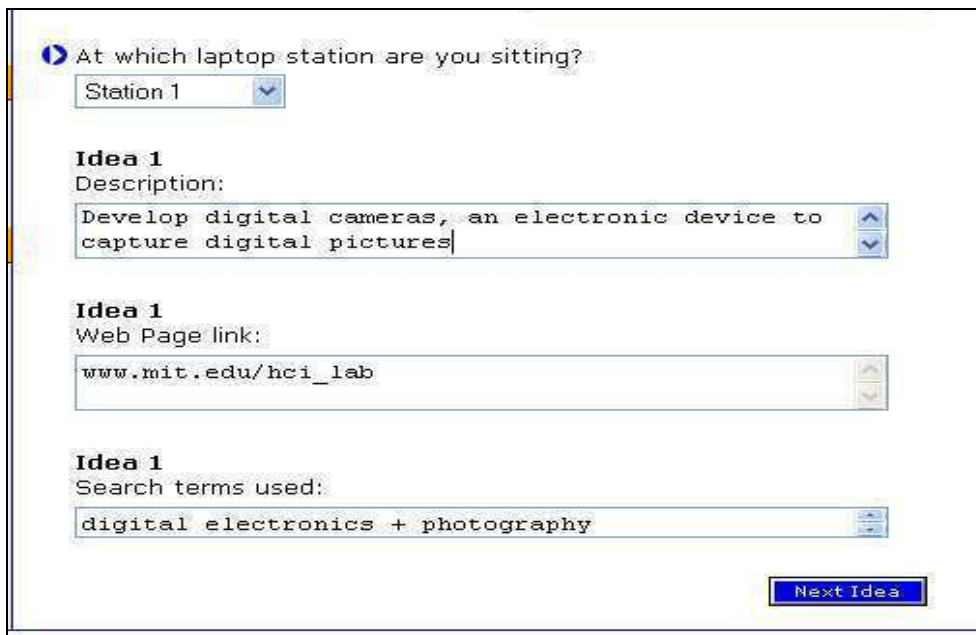
LAB EXPERIMENT – TASK

You work for a company that designs, develops and sells MP3 players. Thus, your company knows and understands MP3 technology and the market for MP3 players very well. They would like to find **novel** ways of applying their MP3 expertise – for example, to develop new services, new products, or new markets.

As a **practical example**, in the 90's Sony applied its extensive knowledge of electronics and digital devices to create a new product and market – digital cameras.

Please use the search tool provided to search the Web for new potential applications of your company's MP3 knowledge. Focus your efforts on using your Web-search results, as provided by the tool, to generate these ideas. Feel free to reuse, modify or create ideas based on what you find. Your objective should be to **generate 4 to 5 ideas** that you think are novel ways of applying your company's MP3 expertise, for example, to develop new services, new products, or new markets. The quality of the idea is more important than quantity.

As you use the tool to search the Web, please **document your ideas** in the Web form ("Ideas") open on your desktop by recording the following: 1) the **idea**, 2) the **webpage** from which you generated or found the idea, and 3) the **search terms** you were using at the time. For example, an entry would look like the following:



The screenshot shows a web form titled "At which laptop station are you sitting?" with a dropdown menu set to "Station 1". Below this, there are three sections, each labeled "Idea 1":

- Idea 1 Description:** A text box containing "Develop digital cameras, an electronic device to capture digital pictures".
- Idea 1 Web Page link:** A text box containing "www.mit.edu/hci_lab".
- Idea 1 Search terms used:** A text box containing "digital electronics + photography".

A "Next Idea" button is located at the bottom right of the form.

Ideas will be evaluated on how **novel** and how **beneficial** the idea is to the company. After you have completed this task, you will be asked to rate and rank your ideas and asked a series of questions about your experience.

APPENDIX H PRELIMINARY ANALYSIS

Table H.1 – Missing Value Analysis

	N	Mean	Std. Deviation	Missing		No. of Extremes(a,b)		Case	Id
				Count	Percent	Low	High		
BROAD3	99	4.72	1.34	0	0	10	6		
BROAD4	98	4.79	1.334	1	1	2	0	52	65
BROAD5	99	5.11	1.421	0	0	1	0		
DEEP2	99	4.63	1.225	0	0	7	3		
DEEP3	99	4.88	1.372	0	0	1	0		
DEEP4	99	4.72	1.371	0	0	1	0		
NOVEL1	98	4.7	1.452	1	1	1	0	93	107
NOVEL2	99	4.43	1.526	0	0	0	0		
NOVEL3	99	4.82	1.445	0	0	3	0		
NOVEL4	99	5	1.491	0	0	2	0		
FORAGE1	99	4.85	1.494	0	0	4	0		
FORAGE2	99	5.02	1.37	0	0	16	0		
STERM1	99	3.9	1.705	0	0	0	0		
STERM2	99	4.48	1.473	0	0	0	0		
INTUIT1	99	4.94	1.26	0	0	0	0		
INTUIT3	98	5.03	1.373	1	1	14	0	54	67
INTUIT4	99	5.03	1.281	0	0	13	0		
INTERP1	98	3.98	1.485	1	1	0	0	99	113
INTERP2	99	4.56	1.465	0	0	0	0		
INTERP3	99	4.95	1.366	0	0	2	0		
TUN2	99	4.78	1.208	0	0	1	0		
TUN3	99	5.01	1.129	0	0	9	0		
TUN4	99	5.11	1.186	0	0	9	0		
TUN5	99	4.53	1.164	0	0	4	2		
MMM1	99	4.79	1.18	0	0	1	0		
MMM3	99	4.39	1.185	0	0	6	2		
MMM4	99	5.09	0.893	0	0	5	0		
MMM5	98	4.84	1.052	1	1	3	3	14	74
MMB3	99	3.52	1.402	0	0	0	0		
MMB4	99	4.2	1.428	0	0	0	0		
MMB6	99	3.72	1.422	0	0	0	0		
MMB7	99	4.19	1.412	0	0	0	0		
LEVEL11	99	5.59	1.325	0	0	5	0		
LEVEL12	97	4.97	1.38	2	2	17	0	8, 25	23, 91
LEVEL21	99	5.3	1.515	0	0	13	0		
LEVEL22	98	5.17	1.293	1	1	13	0	8	23
LEVEL31	99	4.79	1.431	0	0	1	0		
LEVEL32	98	5.06	1.25	1	1	1	0	8	23

	N	Mean	Std. Deviation	Missing		No. of Extremes(a,b)		Case	Id
TRAIN1	99	6.52	0.8	0	0	2	0		
TRAIN3	99	2.82	1.792	0	0	0	0		
CLEAR1	99	6.32	0.843	0	0	4	0		
CLEAR2	99	6.44	0.823	0	0	3	0		
SAT2	99	5.33	1.552	0	0	13	0		
SAT3	99	3.09	1.604	0	0	0	2		
SAT4	99	5.42	1.779	0	0	0	0		
PIIT1	99	4.85	1.6	0	0	0	0		
PIIT2	99	3.82	1.734	0	0	0	0		
PIIT3	99	3.39	1.795	0	0	0	0		
PIIT4	98	4.6	1.565	1	1	0	0	89	97
MP32	99	5	1.34	0	0	15	0		
MP31	99	5.41	1.317	0	0	7	0		
GGL1	99	6.31	0.829	0	0	3	0		
GGL2	99	6.62	0.65	0	0	1	0		
VVS1	99	1.13	0.368	0	0	.	.		
VVS2	99	1.02	0.141	0	0	.	.		
GCSE1	99	6.13	2.136	0	0	0	0		
GCSE2	99	5.3	2.243	0	0	0	0		
GCSE3	99	8.12	1.605	0	0	3	0		
GCSE4	99	8.15	1.554	0	0	1	0		
GCSE5	99	7.04	1.958	0	0	4	0		
GCSE6	99	8.51	1.343	0	0	3	0		
WSE11	99	6.04	1.068	0	0	1	0		
WSE21	99	6.01	1.093	0	0	1	0		
WSE31	99	5.79	1.189	0	0	1	0		
WSE41	99	5.75	1.223	0	0	1	0		
GENDER2	99	1.5455	0.50046	0	0	0	0		
YEAR2	99	2.4848	1.27266	0	0	0	0		
AGE1	99	19.68	2.49	0	0	1	0		
a. Number of cases outside the range (Q1 - 1.5*IQR, Q3 + 1.5*IQR). b. indicates that the inter-quartile range (IQR) is zero.									

Table H.2 – Univariate Outlier Analysis – Pre-transformation

Variable	Group/Tool	Value	Case	ID	COMMENT ¹
BROAD5	G	-3.19179	53	66	CLOSE
CLEAR1	A	-3.16255	8	23	CLOSE
CLEAR1	V	-3.85848	76	41	
CLEAR2	G	-3.28052	34	16	CLOSE
CLEAR2	V	-4.29689	76	41	
FORAGE1	G	-3.70131	36	18	
FORAGE2	G	-3.68597	36	18	
GCSE4	V	-3.32292	73	37	
GCSE6	V	-3.0073	75	39	CLOSE
GGL1	A	-3.29104	25	91	
GGL2	A	-3.50191	25	91	
GGL2	G	-3.27185	38	26	CLOSE
INTUIT4	V	-3.16453	65	29	CLOSE
LEVEL11	G	-3.13015	36	18	CLOSE
LEVEL11	V	-3.33823	76	41	
LEVEL32	V	-3.07545	79	44	CLOSE
MMM4	A	-3.10455	9	24	CLOSE
MMM4	V	-3.6799	76	41	
MMM5	V	-3.51107	76	41	
MP32	V	-3.62866	65	29	
SAT2	G	-3.11546	36	18	CLOSE
SAT2	G	-3.11546	34	16	CLOSE
SAT2	V	-3.67635	76	41	
TRAIN1	A	-4.09065	24	90	
TRAIN1	V	-4.03993	76	41	
VVS1	A	4.52558	12	72	
WSE1	A	-4.13947	32	103	
WSE2	V	-4.17557	91	105	
WSE3	V	-4.127	81	46	

¹ Variables that had standardized values greater than 3.00 but less than 3.29 were noted as approaching outlier status - identified as "CLOSE" in the comment column.

Table H.3 – Normality Analysis: Athens²

Dependent Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
BROAD3	-0.826	0.409	-2.01956	-0.028	0.798	-0.03509
BROAD4	-1.145	0.409	-2.79951	0.309	0.798	0.387218
NOVEL3	-0.82	0.409	-2.00489	0.383	0.798	0.47995
INTUIT3	-0.866	0.409	-2.11736	0.287	0.798	0.359649
TUN2	-1.063	0.409	-2.59902	1.522	0.798	1.907268
TUN3	-1.081	0.409	-2.64303	0.965	0.798	1.209273
TUN4	-0.822	0.409	-2.00978	0.974	0.798	1.220551
MMM4	-1.249	0.409	-3.05379	2.156	0.798	2.701754
Covariate Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
LEVEL11	-1.12	0.409	-2.73839	0.701	0.798	0.878446
LEVEL21	-1.244	0.409	-3.04156	1.073	0.798	1.344612
TRAIN1	-2.225	0.409	-5.4401	7.777	0.798	9.745614
TRAIN3	1.129	0.409	2.760391	0.784	0.798	0.982456
CLEAR1	-1.063	0.409	-2.59902	1.702	0.798	2.132832
GGL1	-1.569	0.409	-3.83619	2.475	0.798	3.101504
GGL2	-2.018	0.409	-4.93399	3.594	0.798	4.503759
VVS1	3.69	0.409	9.022005	14.033	0.798	17.58521
VVS2	5.745	0.409	14.04645	33	0.798	41.35338
GCSE3	-1.266	0.409	-3.09535	1.336	0.798	1.674185
WSE11	-2.523	0.409	-6.1687	8.637	0.798	10.82331
WSE31	-0.839	0.409	-2.05134	0.27	0.798	0.338346
GENDER2	-0.594	0.409	-1.45232	-1.757	0.798	-2.20175
YEAR2	-0.069	0.409	-0.1687	-1.795	0.798	-2.24937

² **Bolded** cells represent negative skewness or kurtosis. Grey highlighting represents positive skewness or kurtosis.

Table H.4 – Normality Analysis: Google

Dependent Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
BROAD5	-0.942	0.421	-2.23753	2.116	0.821	2.577345
FORAGE1	-1.665	0.421	-3.95487	5.297	0.821	6.451888
FORAGE2	-1.67	0.421	-3.96675	5.166	0.821	6.292326
INTUIT3	-1.218	0.427	-2.85246	1.273	0.833	1.528211
INTERP3	-0.897	0.421	-2.13064	0.727	0.821	0.885505
TUN2	-0.844	0.421	-2.00475	0.955	0.821	1.163216
Covariate Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
LEVEL11	-0.936	0.421	-2.22328	1.51	0.821	1.83922
LEVEL12	-1.277	0.421	-3.03325	1.946	0.821	2.37028
TRAIN1	-1.523	0.421	-3.61758	1.122	0.821	1.366626
CLEAR1	-1.091	0.421	-2.59145	0.32	0.821	0.389769
CLEAR2	-1.743	0.421	-4.14014	3.023	0.821	3.682095
SAT2	-1.993	0.421	-4.73397	4.635	0.821	5.645554
SAT4	-1.825	0.421	-4.33492	2.216	0.821	2.699147
MP31	-1.016	0.421	-2.4133	0.62	0.821	0.755177
GGL1	-0.919	0.421	-2.1829	0.393	0.821	0.478685
GGL2	-1.703	0.421	-4.04513	2.766	0.821	3.369062
VVS1	2.327	0.421	5.527316	3.648	0.821	4.443362
GCSE3	-0.853	0.421	-2.02613	0.458	0.821	0.557856
GCSE4	-0.975	0.421	-2.31591	0.14	0.821	0.170524
GCSE6	-0.867	0.421	-2.05938	0.54	0.821	0.657734
WSE31	-0.955	0.421	-2.26841	0.821	0.821	1
GENDER2	-0.068	0.421	-0.16152	-2.138	0.821	-2.60414

Table H.5 – Normality Analysis: Vivisimo

Dependent Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
BROAD3	-0.97	0.398	-2.43719	0.561	0.778	0.72108
DEEP2	-0.831	0.398	-2.08794	-0.005	0.778	-0.00643
FORAGE1	-0.912	0.398	-2.29146	0.49	0.778	0.62982
FORAGE2	-0.77	0.398	-1.93467	1.785	0.778	2.294344
INTUIT4	-1.267	0.398	-3.18342	2.515	0.778	3.232648
TUN2	-0.784	0.398	-1.96985	0.153	0.778	0.196658
MMM4	-1.587	0.398	-3.98744	4.008	0.778	5.151671
MMM5	-1.16	0.398	-2.91457	3.153	0.778	4.052699
MP32	-1.657	0.398	-4.16332	4.296	0.778	5.521851
Covariate Variables						
Variable	Skewness	S.E. Skew	Z	Kurtosis	S.E. Kurt	Z
LEVEL11	-1.225	0.398	-3.07789	2.954	0.778	3.796915
LEVEL21	-1.125	0.398	-2.82663	0.534	0.778	0.686375
LEVEL32	-0.844	0.398	-2.1206	1.634	0.778	2.100257
TRAIN1	-2.532	0.398	-6.36181	7.196	0.778	9.249357
TRAIN3	0.79	0.398	1.984925	-0.636	0.778	-0.81748
CLEAR1	-2.002	0.398	-5.03015	5.332	0.778	6.85347
CLEAR2	-2.614	0.398	-6.56784	9.248	0.778	11.88689
SAT2	-1.73	0.398	-4.34673	4.423	0.778	5.68509
SAT4	-1.042	0.398	-2.61809	-0.211	0.778	-0.27121
GGL2	-1.28	0.398	-3.21608	0.543	0.778	0.697943
VVS1	2.134	0.398	5.361809	2.705	0.778	3.476864
VVS2	5.916	0.398	14.86432	35	0.778	44.98715
GCSE4	-0.897	0.398	-2.25377	1.951	0.778	2.507712
GCSE6	-1.033	0.398	-2.59548	0.933	0.778	1.199229
WSE21	-2.274	0.398	-5.71357	7.943	0.778	10.20951
WSE31	-2.239	0.398	-5.62563	7.778	0.778	9.997429
WSE41	-0.97	0.398	-2.43719	0.9	0.778	1.156812
GENDER2	0.06	0.398	0.150754	-2.121	0.778	-2.72622

Table H.6 – Transformation Summary

Variable	Transformed?	Type?	Comments
BROAD3	TBROAD3	SQRT(8-X)	
BROAD4	TBROAD4	SQRT(8-X)	Still some minor skewness (now positive) for Athens.
BROAD5	TBROAD5	SQRT(8-X)	
DEEP2	TDEEP2	SQRT(8-X)	
DEEP3			
DEEP4			
NOVEL1			
NOVEL2			
NOVEL3	TNOVEL3	SQRT(8-X)	
NOVEL4			
FORAGE1a	TFORAGE1a	LG10(8-X)	
FORAGE2	TFORAGE2	SQRT(8-X)	Still some positive kurtosis in Google.
STERM1			
STERM2			
INTUIT1			
INTUIT3	TINTUIT3	SQRT(8-X)	
INTUIT4	TINTUIT4	SQRT(8-X)	
INTERP1			
INTERP2			
INTERP3	TINTERP3	SQRT(8-X)	
TUN2	TTUN2	SQRT(8-X)	
TUN3	TTUN3	SQRT(8-X)	
TUN4	TTUN4	SQRT(8-X)	
TUN5			
MMM1			
MMM3			
MMM4a	TMMM4a	LG10(8-X)	
MMM5	TMMM5	SQRT(8-X)	
MMB3			
MMB4			
MMB6			
MMB7			
LEVEL11	TLEVEL11	SQRT(8-X)	
LEVEL12	TLEVEL12	SQRT(8-X)	
LEVEL21	TLEVEL21	SQRT(8-X)	
LEVEL22			
LEVEL31			
LEVEL32	TLEVEL32	SQRT(8-X)	
TRAIN1a	TTRAIN1a	LG10(8-X)	Still positive skewness in Google and Vivisimo, and positive kurtosis in Vivisimo. However, normality is better than after the first transformation.
TRAIN3	TTRAIN3	SQRT(X)	
CLEAR1a	TCLEAR1a	LG10(8-X)	Still some positive skewness in Vivisimo. However, normality is better than after the first

Variable	Transformed?	Type?	Comments
			transformation.
CLEAR2a	TCLEAR2a	LG10(8-X)	Still positive skewness in Google and Vivisimo. However, normality is better than after the first transformation.
SAT2a	TSAT2a	LG10(8-X)	
SAT3			
SAT4a	TSAT4a	LG10(8-X)	Still positive skewness in Google. However, normality is better than after the first transformation.
PIIT1			
PIIT2			
PIIT3			
PIIT4			
MP32a	TMP32a	LG10(8-X)	
MP31	TMP31	SQRT(8-X)	
GGL1	TGGL1	SQRT(8-X)	Still positive skewness in Athens
GGL2a	TGGL2a	LG10(8-X)	Still positive skewness in Athens. However, normality is better than after the first transformation.
VVS1a	TVVS1a	LG10(8-X)	Still positive skewness and kurtosis in Athens, Google and Vivisimo. However, normality is better than after the first transformation.
GCSE1			
GCSE2			
GCSE3	TGCSE3	SQRT(8-X)	
GCSE4	TGCSE4	SQRT(8-X)	
GCSE5			
GCSE6	TGCSE6	SQRT(8-X)	
WSE11a	WSE11a	LG10(8-X)	
WSE21a	WSE21a	LG10(8-X)	Created minor negative kurtosis in Athens. However, resolved high positive skewness and kurtosis in Vivisimo.
WSE31a	WSE31a	LG10(8-X)	
WSE41	WSE41	SQRT(8-X)	

Table H.7 – Multivariate Outlier Analysis – Post-transformation

TOOL			Case Number	Statistic
A	Mahal. Distance	1	23	31.030
		2	2	31.030
		3	19	31.030
		4	12	31.030
		5	14	31.030
		6	10	31.030
		7	29	31.030
		8	25	31.030
		9	20	31.030
		10	17	31.030
G	Mahal. Distance	1	61	29.032
		2	49	29.032
		3	63	29.032
		4	42	29.032
		5	58	29.032
		6	34	29.032
		7	43	29.032
		8	50	29.032
		9	54	29.032
		10	53	29.032
V	Mahal. Distance	1	72	33.006
		2	65	32.928
		3	68	32.832
		4	71	32.799
		5	73	32.719
		6	88	32.691
		7	91	32.685
		8	89	32.612
		9	69	32.605
		10	92	32.574

APPENDIX I CORRELATIONS

Table I.1 – Correlation Matrix: Dependent Variables

		TDEEP2	DEEP3	DEEP4	TBROAD3	TBROAD4	TBROAD5	NOVEL1	NOVEL2	TNOVEL3	NOVEL4
TDEEP2		1									
	Sig.										
DEEP3		-0.4211***	1								
	Sig.	0.0000									
DEEP4		-0.5821***	0.4590***	1							
	Sig.	0.0000	0.0000								
TBROAD3		0.5365***	-0.5319***	-0.5794***	1						
	Sig.	0.0000	0.0000	0.0000							
TBROAD4		0.1424	-0.5637***	-0.2650**	0.3317***	1					
	Sig.	0.1597	0.0000	0.0080	0.0008						
TBROAD5		0.3068**	-0.8079***	-0.4528***	0.4240***	0.5255***	1				
	Sig.	0.0020	0.0000	0.0000	0.0000	0.0000					
NOVEL1		-0.2209*	0.5737***	0.4056***	-0.4429***	-0.3271***	-0.5904***	1			
	Sig.	0.0280	0.0000	0.0000	0.0000	0.0010	0.0000				
NOVEL2		-0.1589	0.3323***	0.2349*	-0.2358*	-0.5807***	-0.3932***	0.3526***	1		
	Sig.	0.1162	0.0008	0.0193	0.0188	0.0000	0.0001	0.0003			
TNOVEL3		0.1910	-0.5640***	-0.4426***	0.3869***	0.3292***	0.5184***	-0.6297***	-0.4316***	1	
	Sig.	0.0582	0.0000	0.0000	0.0001	0.0009	0.0000	0.0000	0.0000		
NOVEL4		-0.2263*	0.5484***	0.2496*	-0.3863***	-0.4188***	-0.4506***	0.4102***	0.3003**	-0.4020***	1
	Sig.	0.0243	0.0000	0.0127	0.0001	0.0000	0.0000	0.0000	0.0025	0.0000	
TFORAGE1a		0.3937***	-0.4840***	-0.4887***	0.5448***	0.3179**	0.4815***	-0.4472***	-0.1699	0.3821***	-0.3855***
	Sig.	0.0001	0.0000	0.0000	0.0000	0.0013	0.0000	0.0000	0.0928	0.0001	0.0001
TFORAGE2		0.1445	-0.3815***	-0.3275***	0.3183**	0.1923	0.3745***	-0.2626**	-0.0774	0.4044***	-0.3653***
	Sig.	0.1536	0.0001	0.0009	0.0013	0.0566	0.0001	0.0086	0.4464	0.0000	0.0002
STERM1		-0.0248	-0.0227	-0.1040	0.0296	-0.0036	0.0330	-0.0471	0.0366	-0.0292	-0.1645
	Sig.	0.8077	0.8233	0.3055	0.7713	0.9715	0.7460	0.6431	0.7189	0.7743	0.1037
STERM2		-0.1411	0.1606	0.1949	-0.2859**	-0.0305	-0.1337	0.3397***	0.2094*	-0.3523***	-0.0139
	Sig.	0.1636	0.1123	0.0532	0.0041	0.7644	0.1870	0.0006	0.0375	0.0003	0.8911
INTUIT1		-0.2551*	0.4560***	0.3681***	-0.4292***	-0.4316***	-0.4604***	0.4184***	0.2738**	-0.4455***	0.5863***
	Sig.	0.0108	0.0000	0.0002	0.0000	0.0000	0.0000	0.0000	0.0061	0.0000	0.0000
TINTUIT3		0.1227	-0.4034***	-0.1951	0.3218**	0.3598***	0.3453***	-0.4003***	-0.4404***	0.4813***	-0.4095***
	Sig.	0.2264	0.0000	0.0529	0.0012	0.0003	0.0005	0.0000	0.0000	0.0000	0.0000
TINTUIT4		0.1582	-0.5544***	-0.3208**	0.4056***	0.3944***	0.6139***	-0.4929***	-0.3701***	0.5037***	-0.4996***
	Sig.	0.1177	0.0000	0.0012	0.0000	0.0001	0.0000	0.0000	0.0002	0.0000	0.0000

		TDEEP2	DEEP3	DEEP4	TBROAD3	TBROAD4	TBROAD5	NOVEL1	NOVEL2	TNOVEL3	NOVEL4
INTERP1		0.0641	0.1533	0.1320	-0.1486	-0.2697**	-0.1567	0.2236*	0.2521*	-0.2791**	0.1374
	Sig.	0.5286	0.1297	0.1928	0.1422	0.0069	0.1213	0.0261	0.0118	0.0051	0.1751
INTERP2		-0.2504*	0.3992***	0.4194***	-0.4077***	-0.2918**	-0.3957***	0.4287***	0.0735	-0.3220**	0.2988**
	Sig.	0.0124	0.0000	0.0000	0.0000	0.0034	0.0001	0.0000	0.4697	0.0012	0.0027
TINTERP3		0.0655	-0.0036	-0.1772	0.1517	-0.0178	0.1654	-0.1623	-0.0993	0.1409	-0.1840
	Sig.	0.5192	0.9717	0.0793	0.1340	0.8609	0.1019	0.1085	0.3282	0.1643	0.0683
MMM1		-0.4313***	0.3493***	0.4671***	-0.5199***	-0.1644	-0.2760**	0.2165*	0.1026	-0.3366***	0.1681
	Sig.	0.0000	0.0004	0.0000	0.0000	0.1040	0.0057	0.0314	0.3121	0.0007	0.0963
MMM3		-0.1309	0.2179*	0.1887	-0.3238**	-0.1819	-0.1712	0.2479*	0.1752	-0.2339*	0.0750
	Sig.	0.1964	0.0303	0.0615	0.0011	0.0715	0.0902	0.0134	0.0828	0.0198	0.4603
TMMM4a		0.2375*	-0.3855***	-0.3514***	0.4879***	0.2634**	0.3474***	-0.2815**	-0.2177	0.3729***	-0.3239**
	Sig.	0.0179	0.0001	0.0004	0.0000	0.0084	0.0004	0.0048	0.0305	0.0001	0.0011
TMMM5		0.3491***	-0.4194***	-0.4522***	0.5195***	0.2212*	0.3313***	-0.3874***	-0.2468*	0.4877***	-0.2363*
	Sig.	0.0004	0.0000	0.0000	0.0000	0.0278	0.0008	0.0001	0.0138	0.0000	0.0185
TTUN2		0.3493***	-0.4962***	-0.3648***	0.4554***	0.2996**	0.5555***	-0.4936***	-0.2512*	0.4113***	-0.3926***
	Sig.	0.0004	0.0000	0.0002	0.0000	0.0026	0.0000	0.0000	0.0121	0.0000	0.0001
TTUN3		0.2476*	-0.5787***	-0.4157***	0.3697***	0.2682**	0.6454***	-0.5290***	-0.2491*	0.4296***	-0.3476***
	Sig.	0.0135	0.0000	0.0000	0.0002	0.0073	0.0000	0.0000	0.0129	0.0000	0.0004
TTUN4		0.2755**	-0.6189***	-0.3679***	0.4381***	0.3332***	0.6435***	-0.5899***	-0.3018**	0.5420***	-0.3834***
	Sig.	0.0058	0.0000	0.0002	0.0000	0.0008	0.0000	0.0000	0.0024	0.0000	0.0001
TUN5		-0.3701***	0.5515***	0.5228***	-0.4597***	-0.3865***	-0.5524***	0.4738***	0.3241**	-0.4893***	0.3939***
	Sig.	0.0002	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000	0.0011	0.0000	0.0001
MMB3		0.0994	0.0222	-0.0402	-0.0399	-0.1055	-0.1099	0.1107	0.0803	-0.1265	0.0049
	Sig.	0.3277	0.8275	0.6926	0.6948	0.2986	0.2788	0.2754	0.4294	0.2122	0.9618
MMB4		-0.0675	0.2156*	0.0816	-0.0558	-0.3440***	-0.1963	0.2054*	0.3431***	-0.3266***	0.1868
	Sig.	0.5068	0.0321	0.4220	0.5834	0.0005	0.0515	0.0414	0.0005	0.0010	0.0641
MMB6		-0.0884	0.1078	0.0371	-0.0244	-0.2288*	-0.1614	0.2033*	0.2688**	-0.2161*	-0.0289
	Sig.	0.3841	0.2883	0.7156	0.8102	0.0227	0.1106	0.0436	0.0071	0.0317	0.7766
MMB7		-0.0363	0.2860**	0.1813	-0.0885	-0.3569***	-0.3251**	0.3335***	0.3823***	-0.4093***	0.2374*
	Sig.	0.7212	0.0041	0.0726	0.3837	0.0003	0.0010	0.0007	0.0001	0.0000	0.0180

		TFORAGE1a	TFORAGE2	STERM1	STERM2	INTUIT1	TINTUIT3	TINTUIT4	INTERP1	INTERP2	TINTERP3
TDEEP2	Sig.										
DEEP3	Sig.										
DEEP4	Sig.										
TBROAD3	Sig.										
TBROAD4	Sig.										
TBROAD5	Sig.										
NOVEL1	Sig.										
NOVEL2	Sig.										
TNOVEL3	Sig.										
NOVEL4	Sig.										
TFORAGE1a	Sig.	1									
TFORAGE2	Sig.	0.5382***	1								
STERM1	Sig.	0.0000		1							
STERM2	Sig.	0.1198	0.1359		1						
INTUIT1	Sig.	0.2377	0.1798	0.2065*		1					
TINTUIT3	Sig.	-0.1932	-0.1602	0.0403	0.2468*		1				
TINTUIT4	Sig.	0.0553	0.1132	0.3594	0.0138	-0.4431***		1			
INTERP1	Sig.	-0.5383***	-0.5461***	0.0759	-0.2271*	0.0000	0.5968***		1		
INTERP2	Sig.	0.0000	0.0000	0.4550	0.0238	0.0000	0.0000	0.0000		1	
TINTERP3	Sig.	0.0006	0.0144	0.2142	0.2172	0.0000	0.0000	-0.2738**	-0.2886**		1
	Sig.	0.4283***	0.4831***	0.0492	0.2912**	0.3428***	0.0061	0.0038			
	Sig.	0.0000	0.0000	0.6287	0.0035	0.0005					
	Sig.	-0.2421*	-0.3157**								
	Sig.	0.0158	0.0015								

		TFORAGE1a	TFORAGE2	STERM1	STERM2	INTUIT1	TINTUIT3	TINTUIT4	INTERP1	INTERP2	TINTERP3
INTERP2		-0.5927***	-0.5018***	-0.1447	0.2946**	0.5489***	-0.3278***	-0.3627***	0.4535***	1	
	Sig.	0.0000	0.0000	0.1529	0.0031	0.0000	0.0009	0.0002	0.0000		
TINTERP3		0.3067**	0.2296*	0.0656	-0.0774	-0.3988***	0.1763	0.2609**	-0.3361***	-0.2815**	1
	Sig.	0.0020	0.0223	0.5186	0.4462	0.0000	0.0808	0.0091	0.0007	0.0048	
MMM1		-0.4134***	-0.2192*	-0.0158	0.2064*	0.3960***	-0.3766***	-0.3111**	0.2220*	0.4169***	-0.2532*
	Sig.	0.0000	0.0293	0.8765	0.0404	0.0000	0.0001	0.0017	0.0272	0.0000	0.0115
MMM3		-0.2915**	-0.1920	0.0956	0.4213***	0.3578***	-0.3135**	-0.2810**	0.2814**	0.3898***	-0.1335
	Sig.	0.0034	0.0569	0.3464	0.0000	0.0003	0.0016	0.0048	0.0048	0.0001	0.1876
TMMM4a		0.3635***	0.5159***	0.0596	-0.2122*	-0.5163***	0.3762***	0.5400***	-0.2821**	-0.4044***	0.3215**
	Sig.	0.0002	0.0000	0.5581	0.0350	0.0000	0.0001	0.0000	0.0047	0.0000	0.0012
TMMM5		0.4480***	0.2972**	0.0142	-0.3299***	-0.3577***	0.3813***	0.4374***	-0.1936	-0.4490***	0.1587
	Sig.	0.0000	0.0028	0.8888	0.0009	0.0003	0.0001	0.0000	0.0548	0.0000	0.1166
TTUN2		0.3459***	0.2320*	0.1293	-0.1812	-0.4131***	0.3062**	0.4631***	-0.1044	-0.3423***	0.1635
	Sig.	0.0005	0.0208	0.2023	0.0726	0.0000	0.0021	0.0000	0.3039	0.0005	0.1058
TTUN3		0.4627***	0.3716***	0.1620	0.0094	-0.4169***	0.4370***	0.6177***	-0.1287	-0.3011**	0.1177
	Sig.	0.0000	0.0002	0.1092	0.9264	0.0000	0.0000	0.0000	0.2043	0.0025	0.2459
TTUN4		0.3966***	0.3406***	0.0993	-0.1418	-0.4573***	0.4708***	0.6590***	-0.2172*	-0.3786***	0.1884
	Sig.	0.0000	0.0006	0.3280	0.1615	0.0000	0.0000	0.0000	0.0308	0.0001	0.0619
TUN5		-0.4500***	-0.3452***	-0.0347	0.2428*	0.5090***	-0.5283***	-0.5563***	0.2326*	0.4256***	-0.1868
	Sig.	0.0000	0.0005	0.7331	0.0155	0.0000	0.0000	0.0000	0.0205	0.0000	0.0642
MMB3		-0.0300	-0.0319	0.0647	0.0310	0.1622	-0.1931	-0.1451	0.0419	0.0381	0.1786
	Sig.	0.7681	0.7539	0.5249	0.7608	0.1087	0.0555	0.1520	0.6807	0.7083	0.0769
MMB4		-0.1570	-0.1917	0.0587	0.1421	0.2676**	-0.4512***	-0.3068**	0.1256	0.1262	0.0954
	Sig.	0.1206	0.0573	0.5636	0.1607	0.0074	0.0000	0.0020	0.2156	0.2132	0.3476
MMB6		-0.0906	0.1246	0.2701**	0.1246	0.1042	-0.2069*	-0.0622	0.0624	0.0615	0.0232
	Sig.	0.3725	0.2191	0.0069	0.2192	0.3045	0.0399	0.5406	0.5396	0.5454	0.8197
MMB7		-0.1198	-0.1148	0.0802	0.2491*	0.4367***	-0.4922***	-0.4731***	0.2721**	0.2291*	-0.1962
	Sig.	0.2376	0.2580	0.4301	0.0129	0.0000	0.0000	0.0000	0.0064	0.0226	0.0517

		MMM1	MMM3	TMMM4a	TMMM5	TTUN2	TTUN3	TTUN4	TUN5	MMB3	MMB4	MMB6	MMB7
MMM1		1											
	Sig.												
MMM3		0.4251***	1										
	Sig.	0.0000											
TMMM4a		-0.4616***	-0.4684***	1									
	Sig.	0.0000	0.0000										
TMMM5		-0.5646***	-0.5211***	0.4565***	1								
	Sig.	0.0000	0.0000	0.0000									
TTUN2		-0.4094***	-0.3496***	0.3937***	0.4390***	1							
	Sig.	0.0000	0.0004	0.0001	0.0000								
TTUN3		-0.3558***	-0.2802**	0.4455***	0.3687***	0.4791***	1						
	Sig.	0.0003	0.0050	0.0000	0.0002	0.0000							
TTUN4		-0.4426***	-0.2673**	0.4315***	0.5640***	0.4920***	0.6616***	1					
	Sig.	0.0000	0.0075	0.0000	0.0000	0.0000	0.0000						
TUN5		0.5277***	0.4478***	-0.4195***	-0.5964***	-0.5488***	-0.5690***	-0.6776***	1				
	Sig.	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000					
MMB3		0.0297	-0.0804	0.1125	-0.0326	-0.0744	-0.0840	-0.1787	0.1827	1			
	Sig.	0.7704	0.4290	0.2676	0.7486	0.4640	0.4082	0.0768	0.0703				
MMB4		0.0620	0.1273	-0.0838	-0.1211	-0.1738	-0.1589	-0.1691	0.4082***	0.3652***	1		
	Sig.	0.5422	0.2092	0.4095	0.2326	0.0853	0.1162	0.0943	0.0000	0.0002			
MMB6		0.0369	0.0002	0.1507	-0.0381	-0.1309	-0.0490	-0.0877	0.1771	0.4628***	0.4907***	1	
	Sig.	0.7173	0.9986	0.1364	0.7081	0.1964	0.6298	0.3882	0.0795	0.0000	0.0000		
MMB7		0.0798	0.2288**	-0.1618	-0.2260**	-0.2767**	-0.2387**	-0.3533***	0.5219***	0.3825***	0.6130***	0.4543***	1
	Sig.	0.4325	0.0227	0.1096	0.0245	0.0056	0.0173	0.0003	0.0000	0.0001	0.0000	0.0000	
***		Correlation is significant at the 0.001 level (2-tailed).											
**		Correlation is significant at the 0.01 level (2-tailed).											
*		Correlation is significant at the 0.05 level (2-tailed).											

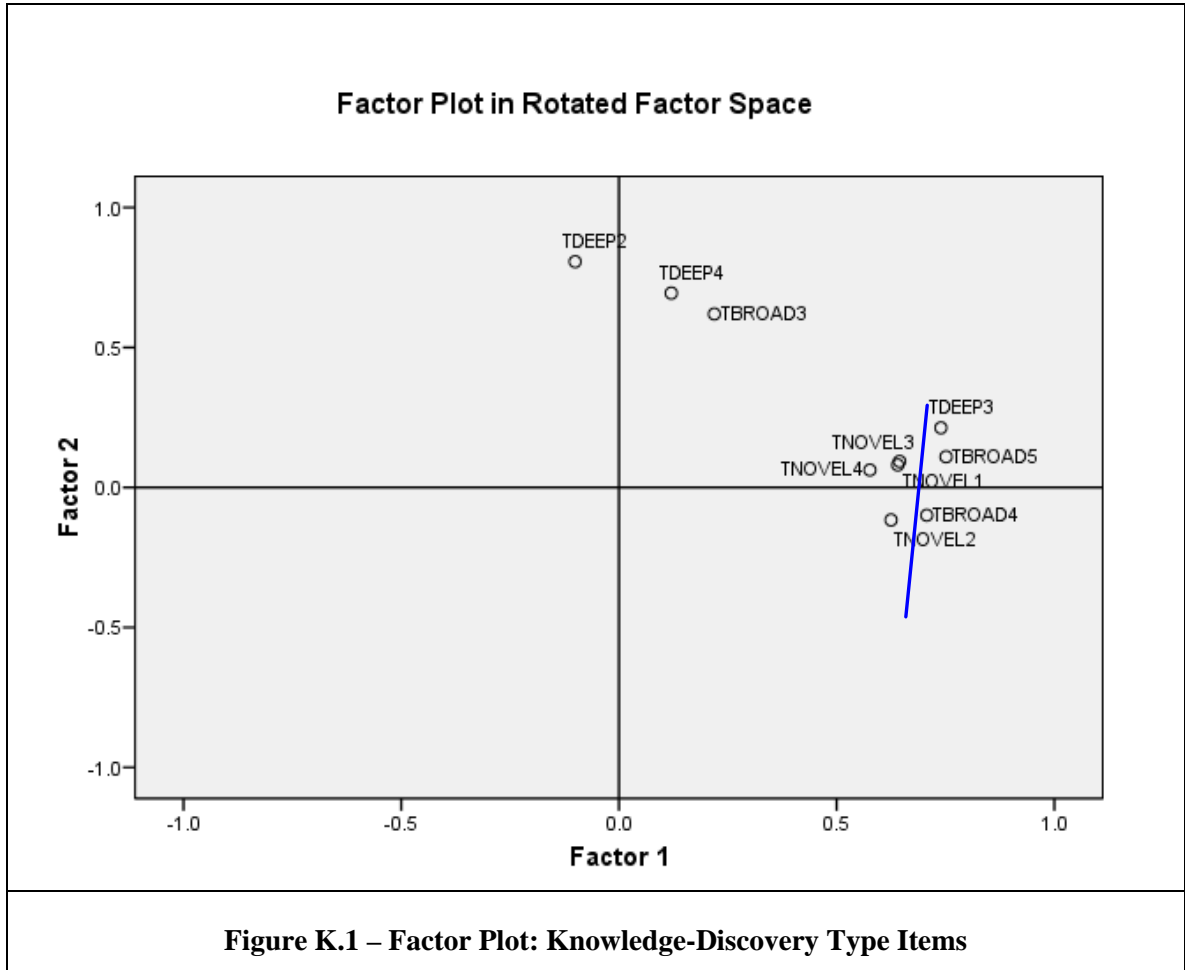
APPENDIX J

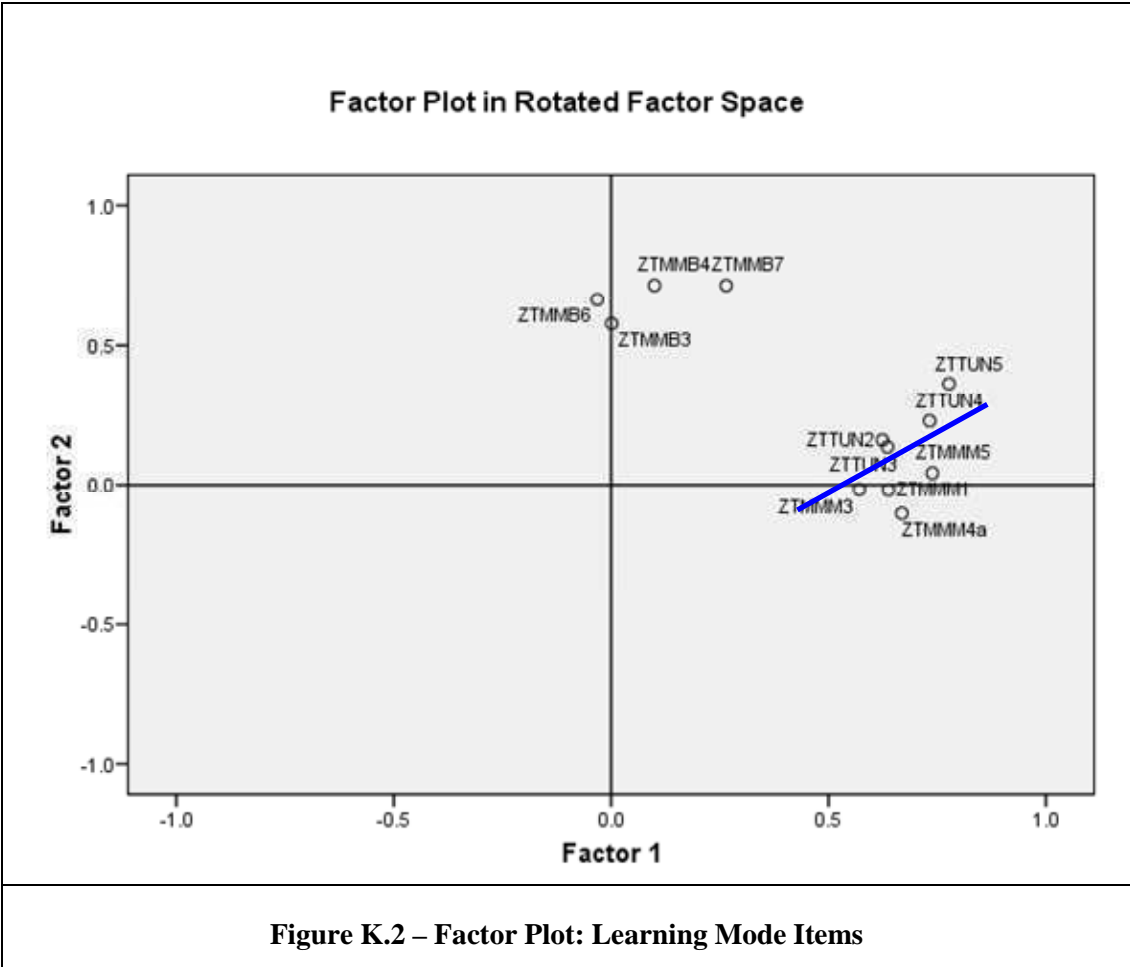
EDITING RULES

The following rules were used when editing the ideas from the lab experiment. The edited ideas were provided to the five independent judges for evaluation.

- Remove duplicate words (for example, “the the” – remove one *the*) or unnecessary and extra words (for example, “into the a computer” – remove *the*)
- Removed unnecessary pluralization.
- Corrected spelling errors (for example, loosing – changed to losing)
- Added words or pluralization as needed (for example, “this idea be targeted to doctors” – added *will*).
- Removed introductions such as “My top ranked is”, “The organization should pursue this idea because”, “The main reason is that”. In some situations, removing this introduction required me to specify what was referred to by “it”. So for example, “the organization should pursue this idea because it will enable it to...” becomes “this will enable the organization to...”. Also added missing qualifiers where needed (for example, “So, if the 30GB Toshiba product ends up being used but none of the movie studios as the base for their HD-DVD discs, Toshiba’s player will be useless.” Added “use 30GB” after Movie studios).
- Removed “I think”.
- Capitalized first letter of words that begin a sentence.
- Generally tried not to refer to the term “idea” in the opening sentence. So instead of “this idea allows”, “this allows”.
- Changed “a mp3” to “an mp3”.
- Added comma before use of the term *but*.
- Removed unnecessary spacing.
- Standardized bullets to a dash plus a space.
- Defining any acronyms that might not be generally known, like DRM.
- Removed exclamation marks.
- Removed reference to Sony (i.e. person assuming name of this company is Sony).

APPENDIX K EXPLORATORY FACTOR ANALYSIS





APPENDIX L

MULTIVARIATE ANALYSIS OF COVARIANCE

Table L.1 – Correlation Matrix: Dependent Variables and Covariates

	TOTDEEP2	TOTBROAD2	TOTNOVEL	TOTSTERM	TOTFORAGE	TOTINTUIT	TOTINTERP	TOTMMM	TOTTUN	TOTMMB	TMP32a
TOTGCSE3	-0.0312	-0.0544	-0.0344	0.0118	-0.1562	-0.1348	-0.1287	-0.0820	0.0916	0.0196	0.1700
Sig.	0.7590	0.5926	0.7351	0.9075	0.1225	0.1833	0.2041	0.4195	0.3672	0.8469	0.0925
TOTWSE2	0.0777	0.0708	-0.0158	-0.0258	0.1091	0.0939	0.1099	0.0870	-0.0198	-0.0148	-0.0248
Sig.	0.4445	0.4865	0.8764	0.8001	0.2824	0.3552	0.2790	0.3917	0.8460	0.8847	0.8072
TOTPIIT2	-0.1825	-0.0034	-0.0082	-0.0708	-0.1861	-0.1805	-0.2201*	-0.2122*	-0.0451	-0.0922	0.0217
Sig.	0.0706	0.9737	0.9355	0.4864	0.0651	0.0738	0.0286	0.0349	0.6573	0.3641	0.8310
TMP31	0.1765	0.0978	0.0996	0.1266	0.1849	0.1454	0.1856	0.2387	-0.0007	-0.0413	-0.0108
Sig.	0.0805	0.3354	0.3265	0.2118	0.0669	0.1509	0.0659	0.0174	0.9945	0.6845	0.9155
TOOLFAMG	0.0324	0.1224	0.0851	0.0432	0.0190	0.0608	-0.0351	0.0313	0.0228	0.0154	0.0340
Sig.	0.7502	0.2276	0.4023	0.6709	0.8519	0.5498	0.7304	0.7584	0.8231	0.8796	0.7382
AGE	0.1544	0.2721**	0.2579**	0.0572	0.2406*	0.1331	0.1223	0.1184	0.2021*	0.0663	0.1834
Sig.	0.1270	0.0064	0.0099	0.5737	0.0165	0.1892	0.2280	0.2431	0.0448	0.5144	0.0692
YEAR	0.1295	0.2175*	0.2569*	0.0241	0.2077*	0.1349	0.1525	0.0900	0.1875	0.0638	0.2172*
Sig.	0.2015	0.0306	0.0103	0.8131	0.0391	0.1831	0.1319	0.3755	0.0631	0.5303	0.0308
GENDER	-0.0034	-0.1973	-0.1165	0.1180	-0.0264	-0.1655	-0.0813	0.0032	-0.1194	-0.1067	-0.0455
Sig.	0.9735	0.0503	0.2508	0.2445	0.7957	0.1017	0.4239	0.9749	0.2393	0.2930	0.6550
***	Correlation is significant at the 0.001 level (2-tailed).										
**	Correlation is significant at the 0.01 level (2-tailed).										
*	Correlation is significant at the 0.05 level (2-tailed).										

Table L.2 – Correlation Matrix: Covariates

		TOTGCSE3	TOTWSE2	TOTPIIT2	TMP31	TOOLFAMG	AGE	YEAR	GENDER
TOTGCSE3		1							
	Sig.								
TOTWSE2		-0.3948***	1						
	Sig.	0.0001							
TOTPIIT2		0.5450***	-0.3631***	1					
	Sig.	0.0000	0.0002						
TMP31		-0.4780***	0.4855***	-0.4175***	1				
	Sig.	0.0000	0.0000	0.0000					
TOOLFAMG		-0.3064**	0.5752***	-0.1775	0.4079***	1			
	Sig.	0.0020	0.0000	0.0788	0.0000				
AGE		0.0051	0.0565	-0.0253	0.1572	0.0144	1		
	Sig.	0.9597	0.5786	0.8039	0.1202	0.8878			
YEAR		0.0503	0.0195	-0.0644	0.0783	-0.0041	0.8847***	1	
	Sig.	0.6207	0.8480	0.5262	0.4411	0.9678	0.0000		
GENDER		-0.2681**	0.1669	-0.3877***	0.1596	-0.1382	0.2005*	0.2694**	1
	Sig.	0.0073	0.0987	0.0001	0.1145	0.1725	0.0466	0.0070	
***		Correlation is significant at the 0.001 level (2-tailed).							
**		Correlation is significant at the 0.01 level (2-tailed).							
*		Correlation is significant at the 0.05 level (2-tailed).							

L.1 MANCOVA – Full Analysis Original Composites

The results of the MANCOVA using the original composite variables are presented below in Tables L.3 through L.6. Table L.3 is also presented in Chapter 7.

Table L.3 – Multivariate Tests of Significance – Full Analysis

Effect	Wilks' Lambda	F	df	p
TMP31	0.897	0.915	10; 80	0.523
TOTPIIT2	0.861	1.289	10; 80	0.251
TOTGCSE3	0.879	1.104	10; 80	0.370
TOTWSE2	0.958	0.348	10; 80	0.964
AGE	0.947	0.447	10; 80	0.919
YEAR	0.937	0.541	10; 80	0.856
GENDER	0.797	2.034	10; 80	0.040
TOOL	0.562	2.670	20; 160	0.000

Table L.4 – Tests of Between-Subjects Effects – Full Analysis

Source	Dependent Variable	SS	df	MS	F	p
Intercept	TOTDEEP2	0.134	1	0.134	1.547	0.217
	TOTBROAD2	0.148	1	0.148	1.460	0.230
	TOTNOVEL	0.316	1	0.316	3.320	0.072
	TOTSTERM	0.835	1	0.835	1.353	0.248
	TOTFORAGE	0.143	1	0.143	0.277	0.600
	TOTINTUIT	0.310	1	0.310	0.485	0.488
	TOTINTERP	0.540	1	0.540	1.068	0.304
	TOTMMM	0.254	1	0.254	0.438	0.510
	TOTTUN	0.284	1	0.284	0.440	0.509
TMP31	TOTMMB	0.004	1	0.004	0.007	0.935
	TOTDEEP2	0.255	1	0.255	2.934	0.090
	TOTBROAD2	0.006	1	0.006	0.060	0.807
	TOTNOVEL	0.099	1	0.099	1.041	0.310
	TOTSTERM	1.541	1	1.541	2.498	0.118
	TOTFORAGE	1.317	1	1.317	2.539	0.115
	TOTINTUIT	0.258	1	0.258	0.403	0.527
	TOTINTERP	0.713	1	0.713	1.407	0.239
	TOTMMM	1.907	1	1.907	3.289	0.073
TOTPIIT2	TOTTUN	0.015	1	0.015	0.024	0.878
	TOTMMB	0.242	1	0.242	0.396	0.531
	TOTDEEP2	0.303	1	0.303	3.486	0.065
	TOTBROAD2	0.006	1	0.006	0.062	0.803
	TOTNOVEL	0.003	1	0.003	0.028	0.868
	TOTSTERM	0.107	1	0.107	0.174	0.678
	TOTFORAGE	1.101	1	1.101	2.123	0.149
	TOTINTUIT	2.069	1	2.069	3.234	0.076

Source	Dependent Variable	SS	df	MS	F	p
	TOTINTERP	1.770	1	1.770	3.496	0.065
	TOTMMM	1.836	1	1.836	3.167	0.079
	TOTTUN	1.430	1	1.430	2.216	0.140
	TOTMMB	1.756	1	1.756	2.877	0.093
TOTGCSE3	TOTDEEP2	0.144	1	0.144	1.663	0.201
	TOTBROAD2	0.061	1	0.061	0.598	0.441
	TOTNOVEL	0.025	1	0.025	0.260	0.612
	TOTSTERM	0.693	1	0.693	1.123	0.292
	TOTFORAGE	0.094	1	0.094	0.182	0.671
	TOTINTUIT	0.233	1	0.233	0.364	0.548
	TOTINTERP	0.010	1	0.010	0.019	0.890
	TOTMMM	0.404	1	0.404	0.697	0.406
	TOTTUN	0.757	1	0.757	1.173	0.282
	TOTMMB	0.073	1	0.073	0.119	0.731
TOTWSE2	TOTDEEP2	0.011	1	0.011	0.130	0.719
	TOTBROAD2	0.028	1	0.028	0.280	0.598
	TOTNOVEL	0.050	1	0.050	0.525	0.471
	TOTSTERM	0.510	1	0.510	0.826	0.366
	TOTFORAGE	0.008	1	0.008	0.015	0.903
	TOTINTUIT	0.001	1	0.001	0.002	0.968
	TOTINTERP	0.002	1	0.002	0.004	0.949
	TOTMMM	0.100	1	0.100	0.173	0.679
	TOTTUN	0.008	1	0.008	0.012	0.914
	TOTMMB	0.004	1	0.004	0.006	0.938
AGE	TOTDEEP2	0.037	1	0.037	0.429	0.514
	TOTBROAD2	0.138	1	0.138	1.357	0.247
	TOTNOVEL	0.006	1	0.006	0.059	0.808
	TOTSTERM	0.224	1	0.224	0.363	0.548
	TOTFORAGE	0.253	1	0.253	0.488	0.487
	TOTINTUIT	0.011	1	0.011	0.017	0.896
	TOTINTERP	0.306	1	0.306	0.604	0.439
	TOTMMM	0.142	1	0.142	0.245	0.622
	TOTTUN	0.340	1	0.340	0.526	0.470
	TOTMMB	0.082	1	0.082	0.134	0.715
YEAR	TOTDEEP2	0.000	1	0.000	0.003	0.958
	TOTBROAD2	0.012	1	0.012	0.119	0.731
	TOTNOVEL	0.148	1	0.148	1.550	0.216
	TOTSTERM	0.314	1	0.314	0.509	0.477
	TOTFORAGE	0.235	1	0.235	0.453	0.503
	TOTINTUIT	0.766	1	0.766	1.198	0.277
	TOTINTERP	1.252	1	1.252	2.473	0.119
	TOTMMM	0.005	1	0.005	0.008	0.929
	TOTTUN	0.090	1	0.090	0.140	0.709
	TOTMMB	0.012	1	0.012	0.020	0.889
GENDER	TOTDEEP2	0.102	1	0.102	1.171	0.282
	TOTBROAD2	1.037	1	1.037	10.225	0.002
	TOTNOVEL	0.423	1	0.423	4.447	0.038

Source	Dependent Variable	SS	df	MS	F	p
	TOTSTERM	0.619	1	0.619	1.004	0.319
	TOTFORAGE	3.120	1	3.120	6.014	0.016
	TOTINTUIT	6.974	1	6.974	10.904	0.001
	TOTINTERP	3.568	1	3.568	7.047	0.009
	TOTMMM	0.646	1	0.646	1.114	0.294
	TOTTUN	2.661	1	2.661	4.122	0.045
	TOTMMB	1.739	1	1.739	2.849	0.095
TOOL	TOTDEEP2	0.647	2	0.324	3.727	0.028
	TOTBROAD2	0.496	2	0.248	2.442	0.093
	TOTNOVEL	0.159	2	0.080	0.836	0.437
	TOTSTERM	1.666	2	0.833	1.350	0.265
	TOTFORAGE	19.733	2	9.866	19.021	0.000
	TOTINTUIT	2.434	2	1.217	1.903	0.155
	TOTINTERP	4.031	2	2.016	3.982	0.022
	TOTMMM	1.341	2	0.670	1.156	0.319
	TOTTUN	2.321	2	1.161	1.798	0.172
	TOTMMB	0.250	2	0.125	0.205	0.815
Error	TOTDEEP2	7.727	89	0.087		
	TOTBROAD2	9.031	89	0.101		
	TOTNOVEL	8.473	89	0.095		
	TOTSTERM	54.912	89	0.617		
	TOTFORAGE	46.166	89	0.519		
	TOTINTUIT	56.925	89	0.640		
	TOTINTERP	45.057	89	0.506		
	TOTMMM	51.603	89	0.580		
	TOTTUN	57.458	89	0.646		
	TOTMMB	54.336	89	0.611		

Table L.5 depicts the adjusted group means for all of the dependent variables. Table L.6 depicts the pairwise comparisons. Note – due to transformations, lower values represent higher scores.

Table L.5 – Adjusted Estimates – Full Analysis

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
TOTDEEP2	A	1.836(a)	.052	1.733	1.939
	G	1.660(a)	.054	1.553	1.768
	V	1.843(a)	.051	1.742	1.945
TOTBROAD2	A	1.810(a)	.056	1.699	1.922
	G	1.648(a)	.058	1.532	1.764
	V	1.668(a)	.055	1.558	1.778
TOTNOVEL	A	1.776(a)	.054	1.668	1.884
	G	1.695(a)	.057	1.583	1.807
	V	1.791(a)	.053	1.684	1.897
TOTSTERM	A	.177(a)	.138	-.098	.452
	G	-.142(a)	.144	-.428	.144
	V	-.041(a)	.136	-.312	.229
TOTFORAGE	A	.524(a)	.127	.272	.776
	G	-.607(a)	.132	-.869	-.344
	V	.043(a)	.125	-.204	.291
TOTINTUIT	A	.185(a)	.141	-.094	.465
	G	-.212(a)	.147	-.503	.079
	V	.013(a)	.139	-.263	.288
TOTINTERP	A	.268(a)	.125	.019	.517
	G	-.237(a)	.130	-.496	.023
	V	-.043(a)	.123	-.288	.202
TOTMMM	A	.154(a)	.134	-.113	.421
	G	-.137(a)	.140	-.415	.140
	V	-.023(a)	.132	-.286	.239
TOTTUN	A	.181(a)	.142	-.100	.462
	G	-.207(a)	.147	-.500	.086
	V	.012(a)	.139	-.264	.289
TOTMMB	A	.054(a)	.138	-.219	.328
	G	-.072(a)	.143	-.356	.213
	V	.012(a)	.135	-.257	.281

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.6 – Pairwise Comparisons – Full Analysis

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
TOTDEEP2	A	G	.176	.075	.064	-.007	.359
		V	-.007	.073	1.000	-.186	.172
	G	A	-.176	.075	.064	-.359	.007
		V	-.183	.076	.052	-.367	.001
	V	A	.007	.073	1.000	-.172	.186
		G	.183	.076	.052	-.001	.367
TOTBROAD2	A	G	.162	.081	.146	-.036	.360
		V	.142	.079	.230	-.052	.336
	G	A	-.162	.081	.146	-.360	.036
		V	-.020	.082	1.000	-.219	.179
	V	A	-.142	.079	.230	-.336	.052
		G	.020	.082	1.000	-.179	.219
TOTNOVEL	A	G	.081	.079	.920	-.111	.273
		V	-.015	.077	1.000	-.202	.173
	G	A	-.081	.079	.920	-.273	.111
		V	-.095	.079	.691	-.289	.098
	V	A	.015	.077	1.000	-.173	.202
		G	.095	.079	.691	-.098	.289
TOTSTERM	A	G	.319	.200	.343	-.169	.808
		V	.218	.196	.803	-.259	.696
	G	A	-.319	.200	.343	-.808	.169
		V	-.101	.201	1.000	-.592	.390
	V	A	-.218	.196	.803	-.696	.259
		G	.101	.201	1.000	-.390	.592
TOTFORAGE	A	G	1.131(*)	.184	.000	.683	1.579
		V	.480(*)	.180	.027	.042	.918
	G	A	-1.131(*)	.184	.000	-1.579	-.683
		V	-.650(*)	.185	.002	-1.101	-.200
	V	A	-.480(*)	.180	.027	-.918	-.042
		G	.650(*)	.185	.002	.200	1.101
TOTINTUIT	A	G	.397	.204	.163	-.100	.895
		V	.173	.199	1.000	-.314	.659
	G	A	-.397	.204	.163	-.895	.100
		V	-.225	.205	.829	-.725	.276
	V	A	-.173	.199	1.000	-.659	.314
		G	.225	.205	.829	-.276	.725
TOTINTERP	A	G	.505(*)	.181	.020	.062	.947
		V	.311	.177	.248	-.121	.744
	G	A	-.505(*)	.181	.020	-.947	-.062
		V	-.193	.182	.876	-.638	.252
	V	A	-.311	.177	.248	-.744	.121
		G	.193	.182	.876	-.252	.638
TOTMMM	A	G	.292	.194	.410	-.182	.765
		V	.177	.190	1.000	-.286	.641
	G	-.292	.194	.410	-.765	.182	

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
TOTTUN	V	V	-.114	.195	1.000	-.590	.362
		A	-.177	.190	1.000	-.641	.286
		G	.114	.195	1.000	-.362	.590
	A	G	.388	.205	.184	-.112	.888
		V	.169	.200	1.000	-.320	.657
		A	-.388	.205	.184	-.888	.112
TOTMMB	V	V	-.219	.206	.870	-.722	.283
		A	-.169	.200	1.000	-.657	.320
		G	.219	.206	.870	-.283	.722
	A	G	.126	.199	1.000	-.360	.612
		V	.042	.195	1.000	-.433	.517
		A	-.126	.199	1.000	-.612	.360
V	V	-.084	.200	1.000	-.572	.405	
	A	-.042	.195	1.000	-.517	.433	
	G	.084	.200	1.000	-.405	.572	

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

L.2 MANCOVA – Item-level Analysis

The results of the item-level MANCOVAs are presented below for knowledge-discovery types, learning processes and learning modes. The untransformed variables were used in this analysis for easier interpretation, which was important for testing theory 3.

L.2.1 Knowledge-Discovery Types

The results of the item-level MANCOVA for knowledge discovery types can be found in Tables L.7 through L.10.

Table L.7 – Multivariate Tests of Significance – Knowledge-Discovery Type

Effect	Wilks' Lambda	F	df	p
TMP31	0.922	0.677	10; 80	0.743
TOTPIIT2	0.906	0.834	10; 80	0.598
TOTGCSE3	0.868	1.219	10; 80	0.292
TOTWSE2	0.862	1.278	10; 80	0.257
AGE	0.893	0.959	10; 80	0.485
YEAR	0.880	1.090	10; 80	0.380
GENDER	0.847	1.443	10; 80	0.177
TOOL	0.766	1.143	20; 160	0.312

Table L.8 – Tests of Between-Subjects Effects – Knowledge-Discovery Type

Source	Dependent Variable	SS	df	MS	F	p	
Intercept	DEEP2	7.089	1	7.089	4.765	0.032	
	DEEP3	10.899	1	10.899	6.407	0.013	
	DEEP4	6.486	1	6.486	3.659	0.059	
	BROAD3	8.021	1	8.021	4.655	0.034	
	BROAD4	2.262	1	2.262	1.382	0.243	
	BROAD5	7.632	1	7.632	4.100	0.046	
	NOVEL1	8.716	1	8.716	4.581	0.035	
	NOVEL2	5.519	1	5.519	2.412	0.124	
	NOVEL3	0.576	1	0.576	0.276	0.601	
	NOVEL4	1.737	1	1.737	0.854	0.358	
	TMP31	DEEP2	1.980	1	1.980	1.331	0.252
	DEEP3	0.000	1	0.000	0.000	0.996	
	DEEP4	0.969	1	0.969	0.547	0.462	
	BROAD3	5.938	1	5.938	3.446	0.067	
	BROAD4	0.051	1	0.051	0.031	0.860	
	BROAD5	0.016	1	0.016	0.008	0.927	
	NOVEL1	0.858	1	0.858	0.451	0.504	
	NOVEL2	2.984	1	2.984	1.304	0.257	
	NOVEL3	0.231	1	0.231	0.111	0.740	
	NOVEL4	0.101	1	0.101	0.050	0.824	
TOTPIIT2	DEEP2	4.127	1	4.127	2.774	0.099	
	DEEP3	0.516	1	0.516	0.304	0.583	
	DEEP4	5.953	1	5.953	3.359	0.070	
	BROAD3	2.063	1	2.063	1.197	0.277	
	BROAD4	1.370	1	1.370	0.837	0.363	
	BROAD5	1.560	1	1.560	0.838	0.362	
	NOVEL1	0.096	1	0.096	0.051	0.822	
	NOVEL2	0.001	1	0.001	0.001	0.981	
	NOVEL3	0.368	1	0.368	0.176	0.676	
	NOVEL4	0.390	1	0.390	0.192	0.663	
	TOTGCSE3	DEEP2	1.741	1	1.741	1.171	0.282
	DEEP3	0.353	1	0.353	0.207	0.650	
	DEEP4	2.038	1	2.038	1.150	0.286	
	BROAD3	3.184	1	3.184	1.848	0.177	
	BROAD4	4.404	1	4.404	2.690	0.104	
	BROAD5	0.072	1	0.072	0.038	0.845	
	NOVEL1	0.333	1	0.333	0.175	0.677	
	NOVEL2	0.327	1	0.327	0.143	0.706	
	NOVEL3	0.392	1	0.392	0.187	0.666	
	NOVEL4	7.536	1	7.536	3.706	0.057	
	TOTWSE2	DEEP2	0.066	1	0.066	0.045	0.833
	DEEP3	0.103	1	0.103	0.061	0.806	
	DEEP4	0.237	1	0.237	0.134	0.715	
	BROAD3	0.549	1	0.549	0.319	0.574	
	BROAD4	0.114	1	0.114	0.070	0.792	

Source	Dependent Variable	SS	df	MS	F	p
	BROAD5	0.518	1	0.518	0.279	0.599
	NOVEL1	4.583	1	4.583	2.409	0.124
	NOVEL2	5.450	1	5.450	2.382	0.126
	NOVEL3	0.007	1	0.007	0.003	0.954
	NOVEL4	1.014	1	1.014	0.499	0.482
AGE	DEEP2	0.797	1	0.797	0.536	0.466
	DEEP3	4.844	1	4.844	2.848	0.095
	DEEP4	0.481	1	0.481	0.271	0.604
	BROAD3	0.252	1	0.252	0.146	0.703
	BROAD4	0.118	1	0.118	0.072	0.789
	BROAD5	1.647	1	1.647	0.885	0.349
	NOVEL1	1.102	1	1.102	0.579	0.449
	NOVEL2	0.118	1	0.118	0.052	0.821
	NOVEL3	1.278	1	1.278	0.612	0.436
	NOVEL4	0.222	1	0.222	0.109	0.742
YEAR	DEEP2	1.462	1	1.462	0.983	0.324
	DEEP3	0.527	1	0.527	0.310	0.579
	DEEP4	0.779	1	0.779	0.439	0.509
	BROAD3	0.139	1	0.139	0.081	0.777
	BROAD4	1.707	1	1.707	1.043	0.310
	BROAD5	0.116	1	0.116	0.062	0.803
	NOVEL1	0.269	1	0.269	0.141	0.708
	NOVEL2	2.559	1	2.559	1.118	0.293
	NOVEL3	6.332	1	6.332	3.031	0.085
	NOVEL4	1.775	1	1.775	0.873	0.353
GENDER	DEEP2	0.286	1	0.286	0.193	0.662
	DEEP3	14.173	1	14.173	8.332	0.005
	DEEP4	3.041	1	3.041	1.716	0.194
	BROAD3	1.476	1	1.476	0.857	0.357
	BROAD4	10.480	1	10.480	6.402	0.013
	BROAD5	11.532	1	11.532	6.195	0.015
	NOVEL1	6.873	1	6.873	3.612	0.061
	NOVEL2	0.168	1	0.168	0.074	0.787
	NOVEL3	5.727	1	5.727	2.742	0.101
	NOVEL4	14.668	1	14.668	7.214	0.009
TOOL	DEEP2	6.486	2	3.243	2.180	0.119
	DEEP3	9.593	2	4.796	2.820	0.065
	DEEP4	9.152	2	4.576	2.582	0.081
	BROAD3	10.922	2	5.461	3.170	0.047
	BROAD4	4.212	2	2.106	1.287	0.281
	BROAD5	11.984	2	5.992	3.219	0.045
	NOVEL1	11.050	2	5.525	2.904	0.060
	NOVEL2	0.773	2	0.387	0.169	0.845
	NOVEL3	4.263	2	2.132	1.020	0.365
	NOVEL4	0.028	2	0.014	0.007	0.993
Error	DEEP2	132.396	89	1.488		
	DEEP3	151.393	89	1.701		

Source	Dependent Variable	SS	df	MS	F	p
	DEEP4	157.743	89	1.772		
	BROAD3	153.338	89	1.723		
	BROAD4	145.692	89	1.637		
	BROAD5	165.669	89	1.861		
	NOVEL1	169.343	89	1.903		
	NOVEL2	203.654	89	2.288		
	NOVEL3	185.927	89	2.089		
	NOVEL4	180.968	89	2.033		

Table L.9 – Adjusted Estimates – Knowledge-Discovery Type

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
DEEP2	A	4.450(a)	.215	4.023	4.877
	G	5.017(a)	.224	4.572	5.461
	V	4.447(a)	.211	4.027	4.867
DEEP3	A	4.437(a)	.230	3.980	4.893
	G	5.166(a)	.239	4.691	5.641
	V	5.041(a)	.226	4.592	5.490
DEEP4	A	4.583(a)	.235	4.117	5.049
	G	5.174(a)	.244	4.689	5.659
	V	4.439(a)	.231	3.980	4.897
BROAD3	A	4.406(a)	.231	3.947	4.866
	G	5.211(a)	.241	4.733	5.689
	V	4.573(a)	.227	4.121	5.025
BROAD4	A	4.553(a)	.225	4.105	5.000
	G	4.725(a)	.235	4.259	5.191
	V	5.057(a)	.222	4.616	5.497
BROAD5	A	4.658(a)	.240	4.180	5.135
	G	5.533(a)	.250	5.036	6.030
	V	5.165(a)	.236	4.695	5.635
NOVEL1	A	4.437(a)	.243	3.954	4.920
	G	5.210(a)	.253	4.707	5.712
	V	4.505(a)	.239	4.030	4.979
NOVEL2	A	4.494(a)	.266	3.965	5.024
	G	4.510(a)	.277	3.959	5.061
	V	4.310(a)	.262	3.790	4.831
NOVEL3	A	4.628(a)	.255	4.122	5.134
	G	5.128(a)	.265	4.601	5.655
	V	4.723(a)	.250	4.225	5.220
NOVEL4	A	5.008(a)	.251	4.508	5.507
	G	4.975(a)	.261	4.455	5.494
	V	5.015(a)	.247	4.524	5.506

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.10 – Pairwise Comparisons – Knowledge-Discovery Type

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
DEEP2	A	G	-.567	.311	.215	-1.326	.192
		V	.003	.304	1.000	-.739	.744
	G	A	.567	.311	.215	-.192	1.326
		V	.570	.313	.215	-.193	1.333
	V	A	-.003	.304	1.000	-.744	.739
		G	-.570	.313	.215	-1.333	.193
DEEP3	A	G	-.729	.332	.093	-1.540	.082
		V	-.604	.325	.199	-1.397	.189
	G	A	.729	.332	.093	-.082	1.540
		V	.125	.334	1.000	-.691	.940
	V	A	.604	.325	.199	-.189	1.397
		G	-.125	.334	1.000	-.940	.691
DEEP4	A	G	-.591	.339	.255	-1.419	.237
		V	.145	.332	1.000	-.665	.954
	G	A	.591	.339	.255	-.237	1.419
		V	.736	.341	.101	-.097	1.569
	V	A	-.145	.332	1.000	-.954	.665
		G	-.736	.341	.101	-1.569	.097
BROAD3	A	G	-.805	.335	.055	-1.621	.012
		V	-.166	.327	1.000	-.965	.632
	G	A	.805	.335	.055	-.012	1.621
		V	.638	.336	.183	-.183	1.459
	V	A	.166	.327	1.000	-.632	.965
		G	-.638	.336	.183	-1.459	.183
BROAD4	A	G	-.173	.326	1.000	-.968	.623
		V	-.504	.319	.352	-1.282	.274
	G	A	.173	.326	1.000	-.623	.968
		V	-.332	.328	.944	-1.132	.469
	V	A	.504	.319	.352	-.274	1.282
		G	.332	.328	.944	-.469	1.132
BROAD5	A	G	-.876(*)	.348	.041	-1.724	-.027
		V	-.507	.340	.417	-1.337	.322
	G	A	.876(*)	.348	.041	.027	1.724
		V	.368	.350	.886	-.485	1.222
	V	A	.507	.340	.417	-.322	1.337
		G	-.368	.350	.886	-1.222	.485
NOVEL1	A	G	-.773	.352	.092	-1.631	.085
		V	-.068	.344	1.000	-.906	.771
	G	A	.773	.352	.092	-.085	1.631
		V	.705	.354	.147	-.158	1.568
	V	A	.068	.344	1.000	-.771	.906
		G	-.705	.354	.147	-1.568	.158
NOVEL2	A	G	-.016	.386	1.000	-.957	.925
		V	.184	.377	1.000	-.736	1.104

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
NOVEL3	G	A	.016	.386	1.000	-.925	.957
		V	.200	.388	1.000	-.746	1.146
	V	A	-.184	.377	1.000	-1.104	.736
		G	-.200	.388	1.000	-1.146	.746
	A	G	-.500	.368	.535	-1.399	.399
		V	-.095	.360	1.000	-.974	.784
NOVEL4	G	A	.500	.368	.535	-.399	1.399
		V	.405	.371	.832	-.499	1.309
	V	A	.095	.360	1.000	-.784	.974
		G	-.405	.371	.832	-1.309	.499
	A	G	.033	.364	1.000	-.854	.920
		V	-.008	.355	1.000	-.875	.860
G	A	V	-.033	.364	1.000	-.920	.854
		V	-.040	.366	1.000	-.932	.852
	V	A	.008	.355	1.000	-.860	.875
		G	.040	.366	1.000	-.852	.932

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

L.2.2 Learning Processes

The results of the item-level MANCOVA for learning processes can be found in Tables

L.11 through L.14.

Table L.11 – Multivariate Tests of Significance – Learning Processes

Effect	Wilks' Lambda	F	df	p
TMP31	0.893	0.955	10; 80	0.489
TOTPIIT2	0.904	0.849	10; 80	0.584
TOTGCSE3	0.866	1.240	10; 80	0.279
TOTWSE2	0.958	0.355	10; 80	0.962
AGE	0.898	0.904	10; 80	0.534
YEAR	0.907	0.816	10; 80	0.614
GENDER	0.827	1.669	10; 80	0.103
TOOL	0.599	2.341	20; 160	0.002

Table L.12 – Tests of Between-Subjects Effects – Learning Processes

Source	Dependent Variable	SS	df	MS	F	p
Intercept	STERM1	2.880	1	2.880	0.951	0.332
	STERM2	14.971	1	14.971	7.028	0.009
	FORAGE1	7.795	1	7.795	4.387	0.039
	FORAGE2	5.131	1	5.131	3.599	0.061
	INTUIT1	1.126	1	1.126	0.840	0.362
	INTUIT3	3.076	1	3.076	1.658	0.201
	INTUIT4	0.073	1	0.073	0.047	0.829
	INTERP1	0.841	1	0.841	0.415	0.521
	INTERP2	0.054	1	0.054	0.029	0.866
	INTERP3	5.256	1	5.256	2.906	0.092
TMP31	STERM1	4.651	1	4.651	1.535	0.219
	STERM2	2.449	1	2.449	1.150	0.286
	FORAGE1	6.460	1	6.460	3.636	0.060
	FORAGE2	0.467	1	0.467	0.328	0.568
	INTUIT1	1.642	1	1.642	1.225	0.271
	INTUIT3	0.036	1	0.036	0.019	0.890
	INTUIT4	0.058	1	0.058	0.037	0.847
	INTERP1	0.363	1	0.363	0.179	0.673
	INTERP2	0.026	1	0.026	0.014	0.906
	INTERP3	4.826	1	4.826	2.668	0.106
TOTPIIT2	STERM1	1.564	1	1.564	0.516	0.474
	STERM2	5.613	1	5.613	2.635	0.108
	FORAGE1	0.891	1	0.891	0.501	0.481
	FORAGE2	4.133	1	4.133	2.899	0.092
	INTUIT1	7.368	1	7.368	5.497	0.021
	INTUIT3	0.893	1	0.893	0.481	0.490
	INTUIT4	3.554	1	3.554	2.285	0.134
	INTERP1	3.522	1	3.522	1.738	0.191
	INTERP2	3.099	1	3.099	1.652	0.202
	INTERP3	4.746	1	4.746	2.624	0.109
TOTGCSE3	STERM1	0.177	1	0.177	0.058	0.809
	STERM2	5.068	1	5.068	2.379	0.126
	FORAGE1	0.011	1	0.011	0.006	0.936
	FORAGE2	0.485	1	0.485	0.340	0.561
	INTUIT1	2.176	1	2.176	1.623	0.206
	INTUIT3	0.001	1	0.001	0.001	0.979
	INTUIT4	0.065	1	0.065	0.042	0.838
	INTERP1	0.030	1	0.030	0.015	0.904
	INTERP2	3.601	1	3.601	1.920	0.169
	INTERP3	2.731	1	2.731	1.510	0.222
TOTWSE2	STERM1	0.000	1	0.000	0.000	0.990
	STERM2	4.758	1	4.758	2.234	0.139
	FORAGE1	1.071	1	1.071	0.603	0.440
	FORAGE2	0.065	1	0.065	0.046	0.831
	INTUIT1	0.230	1	0.230	0.172	0.679

Source	Dependent Variable	SS	df	MS	F	p
	INTUIT3	0.119	1	0.119	0.064	0.801
	INTUIT4	0.038	1	0.038	0.025	0.875
	INTERP1	0.079	1	0.079	0.039	0.844
	INTERP2	0.162	1	0.162	0.086	0.769
	INTERP3	0.190	1	0.190	0.105	0.746
AGE	STERM1	0.192	1	0.192	0.063	0.802
	STERM2	2.713	1	2.713	1.274	0.262
	FORAGE1	0.920	1	0.920	0.518	0.474
	FORAGE2	0.965	1	0.965	0.677	0.413
	INTUIT1	0.035	1	0.035	0.026	0.871
	INTUIT3	0.057	1	0.057	0.031	0.861
	INTUIT4	0.974	1	0.974	0.626	0.431
	INTERP1	3.912	1	3.912	1.930	0.168
	INTERP2	0.295	1	0.295	0.158	0.692
	INTERP3	0.058	1	0.058	0.032	0.859
YEAR	STERM1	0.353	1	0.353	0.117	0.734
	STERM2	1.272	1	1.272	0.597	0.442
	FORAGE1	0.047	1	0.047	0.027	0.871
	FORAGE2	0.314	1	0.314	0.220	0.640
	INTUIT1	0.221	1	0.221	0.165	0.686
	INTUIT3	1.411	1	1.411	0.761	0.385
	INTUIT4	3.600	1	3.600	2.314	0.132
	INTERP1	7.488	1	7.488	3.695	0.058
	INTERP2	2.346	1	2.346	1.251	0.266
	INTERP3	0.178	1	0.178	0.099	0.754
GENDER	STERM1	1.744	1	1.744	0.576	0.450
	STERM2	1.317	1	1.317	0.618	0.434
	FORAGE1	7.073	1	7.073	3.980	0.049
	FORAGE2	7.611	1	7.611	5.339	0.023
	INTUIT1	12.860	1	12.860	9.595	0.003
	INTUIT3	9.513	1	9.513	5.127	0.026
	INTUIT4	13.427	1	13.427	8.632	0.004
	INTERP1	10.788	1	10.788	5.323	0.023
	INTERP2	7.769	1	7.769	4.142	0.045
	INTERP3	5.869	1	5.869	3.245	0.075
TOOL	STERM1	6.508	2	3.254	1.074	0.346
	STERM2	2.772	2	1.386	0.651	0.524
	FORAGE1	43.595	2	21.797	12.267	0.000
	FORAGE2	35.309	2	17.654	12.384	0.000
	INTUIT1	6.825	2	3.413	2.546	0.084
	INTUIT3	2.736	2	1.368	0.737	0.481
	INTUIT4	5.900	2	2.950	1.897	0.156
	INTERP1	15.659	2	7.830	3.863	0.025
	INTERP2	24.628	2	12.314	6.566	0.002
	INTERP3	3.998	2	1.999	1.105	0.336
Error	STERM1	269.626	89	3.030		
	STERM2	189.579	89	2.130		

Source	Dependent Variable	SS	df	MS	F	p
	FORAGE1	158.146	89	1.777		
	FORAGE2	126.877	89	1.426		
	INTUIT1	119.292	89	1.340		
	INTUIT3	165.135	89	1.855		
	INTUIT4	138.434	89	1.555		
	INTERP1	180.374	89	2.027		
	INTERP2	166.911	89	1.875		
	INTERP3	160.991	89	1.809		

Table L.13 – Adjusted Estimates – Learning Processes

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
STERM1	A	3.548(a)	.307	2.939	4.157
	G	4.178(a)	.319	3.544	4.812
	V	3.982(a)	.302	3.383	4.582
STERM2	A	4.261(a)	.257	3.751	4.772
	G	4.679(a)	.268	4.147	5.211
	V	4.524(a)	.253	4.021	5.026
FORAGE1	A	4.084(a)	.235	3.618	4.551
	G	5.762(a)	.244	5.276	6.247
	V	4.760(a)	.231	4.301	5.219
FORAGE2	A	4.256(a)	.210	3.838	4.674
	G	5.765(a)	.219	5.330	6.200
	V	5.081(a)	.207	4.670	5.492
INTUIT1	A	4.667(a)	.204	4.262	5.072
	G	5.319(a)	.212	4.897	5.741
	V	4.860(a)	.201	4.462	5.259
INTUIT3	A	4.805(a)	.240	4.328	5.282
	G	5.215(a)	.250	4.719	5.711
	V	5.084(a)	.236	4.615	5.553
INTUIT4	A	4.719(a)	.220	4.283	5.156
	G	5.337(a)	.229	4.882	5.791
	V	5.052(a)	.216	4.623	5.481
INTERP1	A	3.413(a)	.251	2.914	3.911
	G	4.218(a)	.261	3.699	4.737
	V	4.313(a)	.247	3.823	4.803
INTERP2	A	3.869(a)	.241	3.390	4.348
	G	5.090(a)	.251	4.591	5.589
	V	4.729(a)	.237	4.258	5.201
INTERP3	A	4.915(a)	.237	4.445	5.386
	G	5.236(a)	.247	4.746	5.726
	V	4.728(a)	.233	4.265	5.191

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.14 – Pairwise Comparisons – Learning Processes

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
STERM1	A	G	-.630	.444	.477	-1.713	.452
		V	-.434	.434	.958	-1.493	.624
	G	A	.630	.444	.477	-.452	1.713
		V	.196	.446	1.000	-.893	1.285
	V	A	.434	.434	.958	-.624	1.493
		G	-.196	.446	1.000	-1.285	.893
STERM2	A	G	-.418	.372	.794	-1.325	.490
		V	-.262	.364	1.000	-1.150	.625
	G	A	.418	.372	.794	-.490	1.325
		V	.155	.374	1.000	-.757	1.068
	V	A	.262	.364	1.000	-.625	1.150
		G	-.155	.374	1.000	-1.068	.757
FORAGE1	A	G	-1.678(*)	.340	.000	-2.507	-.848
		V	-.676	.332	.134	-1.487	.134
	G	A	1.678(*)	.340	.000	.848	2.507
		V	1.001(*)	.342	.013	.168	1.835
	V	A	.676	.332	.134	-.134	1.487
		G	-1.001(*)	.342	.013	-1.835	-.168
FORAGE2	A	G	-1.509(*)	.304	.000	-2.252	-.766
		V	-.824(*)	.298	.020	-1.551	-.098
	G	A	1.509(*)	.304	.000	.766	2.252
		V	.685	.306	.083	-.062	1.432
	V	A	.824(*)	.298	.020	.098	1.551
		G	-.685	.306	.083	-1.432	.062
INTUIT1	A	G	-.652	.295	.089	-1.372	.068
		V	-.193	.289	1.000	-.897	.511
	G	A	.652	.295	.089	-.068	1.372
		V	.459	.297	.377	-.265	1.183
	V	A	.193	.289	1.000	-.511	.897
		G	-.459	.297	.377	-1.183	.265
INTUIT3	A	G	-.410	.347	.723	-1.257	.438
		V	-.279	.339	1.000	-1.107	.550
	G	A	.410	.347	.723	-.438	1.257
		V	.131	.349	1.000	-.721	.983
	V	A	.279	.339	1.000	-.550	1.107
		G	-.131	.349	1.000	-.983	.721
INTUIT4	A	G	-.617	.318	.166	-1.393	.158
		V	-.333	.311	.862	-1.091	.426
	G	A	.617	.318	.166	-.158	1.393
		V	.285	.320	1.000	-.496	1.065
	V	A	.333	.311	.862	-.426	1.091
		G	-.285	.320	1.000	-1.065	.496
INTERP1	A	G	-.805	.363	.087	-1.691	.080
		V	-.900(*)	.355	.039	-1.766	-.035
	G	.805	.363	.087	-.080	1.691	

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
INTERP2	V	V	-.095	.365	1.000	-.985	.796
		A	.900(*)	.355	.039	.035	1.766
		G	.095	.365	1.000	-.796	.985
	A	G	-1.221(*)	.349	.002	-2.073	-.369
		V	-.860(*)	.341	.041	-1.693	-.027
		A	1.221(*)	.349	.002	.369	2.073
INTERP3	V	V	.361	.351	.921	-.496	1.217
		A	.860(*)	.341	.041	.027	1.693
		G	-.361	.351	.921	-1.217	.496
	A	G	-.321	.343	1.000	-1.157	.516
		V	.188	.335	1.000	-.630	1.005
		A	.321	.343	1.000	-.516	1.157
V	V	.508	.345	.432	-.333	1.350	
	A	-.188	.335	1.000	-1.005	.630	
	G	-.508	.345	.432	-1.350	.333	

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

L.2.3 Learning Modes

The results of the item-level MANCOVA for learning modes can be found in Tables L.15 through L.18.

Table L.15 – Multivariate Tests of Significance – Learning Modes

Effect	Wilks' Lambda	F	df	p
TMP31	0.904	0.689	12; 78	0.757
TOTPIIT2	0.873	0.944	12; 78	0.509
TOTGCSE3	0.941	0.409	12; 78	0.956
TOTWSE2	0.915	0.606	12; 78	0.831
AGE	0.913	0.618	12; 78	0.821
YEAR	0.951	0.338	12; 78	0.979
GENDER	0.812	1.509	12; 78	0.139
TOOL	0.782	0.850	24; 156	0.669

Table L.16 – Tests of Between-Subjects Effects – Learning Modes

Source	Dependent Variable	SS	df	MS	F	p	
Intercept	MMM1	2.069	1	2.069	1.452	0.231	
	MMM3	5.961	1	5.961	4.160	0.044	
	MMM4	7.820	1	7.820	10.863	0.001	
	MMM5	3.974	1	3.974	3.621	0.060	
	TUN2	4.485	1	4.485	3.194	0.077	
	TUN3	11.585	1	11.585	9.191	0.003	
	TUN4	2.795	1	2.795	2.091	0.152	
TUN5	TUN5	4.686	1	4.686	3.577	0.062	
	MMB3	0.000	1	0.000	0.000	0.995	
	MMB4	2.861	1	2.861	1.429	0.235	
	MMB6	2.442	1	2.442	1.133	0.290	
	MMB7	2.681	1	2.681	1.319	0.254	
	TMP31	MMM1	0.993	1	0.993	0.697	0.406
		MMM3	2.610	1	2.610	1.822	0.181
MMM4		1.498	1	1.498	2.081	0.153	
MMM5		2.207	1	2.207	2.010	0.160	
TUN2		0.271	1	0.271	0.193	0.662	
TUN3		0.385	1	0.385	0.305	0.582	
TUN4		0.391	1	0.391	0.292	0.590	
TUN5	TUN5	0.450	1	0.450	0.343	0.559	
	MMB3	1.011	1	1.011	0.500	0.481	
	MMB4	2.130	1	2.130	1.064	0.305	
	MMB6	0.468	1	0.468	0.217	0.642	
	MMB7	0.010	1	0.010	0.005	0.944	
	TOTPIIT2	MMM1	4.235	1	4.235	2.972	0.088
		MMM3	2.070	1	2.070	1.445	0.233
MMM4		3.278	1	3.278	4.554	0.036	
MMM5		0.627	1	0.627	0.571	0.452	
TUN2		3.554	1	3.554	2.530	0.115	
TUN3		1.370	1	1.370	1.087	0.300	
TUN4		1.219	1	1.219	0.912	0.342	
TUN5	TUN5	2.992	1	2.992	2.284	0.134	
	MMB3	7.835	1	7.835	3.875	0.052	
	MMB4	3.187	1	3.187	1.592	0.210	
	MMB6	1.880	1	1.880	0.873	0.353	
	MMB7	4.572	1	4.572	2.249	0.137	
	TOTGCSE3	MMM1	0.800	1	0.800	0.561	0.456
		MMM3	0.188	1	0.188	0.131	0.718
MMM4		0.078	1	0.078	0.108	0.743	
MMM5		2.332	1	2.332	2.124	0.149	
TUN2		0.309	1	0.309	0.220	0.640	
TUN3		0.679	1	0.679	0.539	0.465	
TUN4		1.789	1	1.789	1.338	0.250	
TUN5	TUN5	4.005	1	4.005	3.058	0.084	
	MMB3	0.158	1	0.158	0.078	0.781	

Source	Dependent Variable	SS	df	MS	F	p
	MMB4	0.122	1	0.122	0.061	0.806
	MMB6	0.110	1	0.110	0.051	0.821
	MMB7	2.248	1	2.248	1.106	0.296
TOTWSE2	MMM1	0.134	1	0.134	0.094	0.760
	MMM3	0.154	1	0.154	0.107	0.744
	MMM4	0.000	1	0.000	0.000	0.997
	MMM5	1.404	1	1.404	1.279	0.261
	TUN2	0.124	1	0.124	0.088	0.767
	TUN3	0.005	1	0.005	0.004	0.952
	TUN4	0.201	1	0.201	0.151	0.699
	TUN5	0.021	1	0.021	0.016	0.899
	MMB3	0.562	1	0.562	0.278	0.599
	MMB4	2.688	1	2.688	1.343	0.250
	MMB6	1.157	1	1.157	0.537	0.466
	MMB7	0.023	1	0.023	0.011	0.916
AGE	MMM1	0.033	1	0.033	0.023	0.879
	MMM3	0.502	1	0.502	0.351	0.555
	MMM4	1.004	1	1.004	1.394	0.241
	MMM5	0.053	1	0.053	0.049	0.826
	TUN2	0.718	1	0.718	0.511	0.476
	TUN3	3.147	1	3.147	2.496	0.118
	TUN4	0.060	1	0.060	0.045	0.833
	TUN5	0.038	1	0.038	0.029	0.865
	MMB3	0.315	1	0.315	0.156	0.694
	MMB4	0.954	1	0.954	0.476	0.492
	MMB6	0.483	1	0.483	0.224	0.637
	MMB7	0.050	1	0.050	0.024	0.876
YEAR	MMM1	0.005	1	0.005	0.004	0.952
	MMM3	0.029	1	0.029	0.020	0.888
	MMM4	0.092	1	0.092	0.128	0.722
	MMM5	0.193	1	0.193	0.176	0.676
	TUN2	0.336	1	0.336	0.240	0.626
	TUN3	0.485	1	0.485	0.385	0.537
	TUN4	0.739	1	0.739	0.553	0.459
	TUN5	0.620	1	0.620	0.473	0.493
	MMB3	0.262	1	0.262	0.129	0.720
	MMB4	0.006	1	0.006	0.003	0.956
	MMB6	0.069	1	0.069	0.032	0.859
	MMB7	0.780	1	0.780	0.384	0.537
GENDER	MMM1	2.357	1	2.357	1.654	0.202
	MMM3	0.352	1	0.352	0.245	0.621
	MMM4	0.645	1	0.645	0.896	0.346
	MMM5	0.663	1	0.663	0.604	0.439
	TUN2	8.804	1	8.804	6.269	0.014
	TUN3	1.168	1	1.168	0.927	0.338
	TUN4	6.446	1	6.446	4.823	0.031
	TUN5	1.703	1	1.703	1.300	0.257

Source	Dependent Variable	SS	df	MS	F	p
	MMB3	2.970	1	2.970	1.469	0.229
	MMB4	10.799	1	10.799	5.395	0.022
	MMB6	0.379	1	0.379	0.176	0.676
	MMB7	6.469	1	6.469	3.183	0.078
TOOL	MMM1	1.232	2	0.616	0.432	0.650
	MMM3	0.605	2	0.302	0.211	0.810
	MMM4	1.804	2	0.902	1.253	0.291
	MMM5	4.145	2	2.073	1.888	0.157
	TUN2	2.364	2	1.182	0.842	0.434
	TUN3	3.506	2	1.753	1.391	0.254
	TUN4	9.479	2	4.739	3.546	0.033
	TUN5	4.780	2	2.390	1.824	0.167
	MMB3	1.344	2	0.672	0.332	0.718
	MMB4	5.988	2	2.994	1.496	0.230
	MMB6	1.761	2	0.880	0.409	0.666
	MMB7	0.024	2	0.012	0.006	0.994
Error	MMM1	126.832	89	1.425		
	MMM3	127.513	89	1.433		
	MMM4	64.071	89	0.720		
	MMM5	97.696	89	1.098		
	TUN2	124.992	89	1.404		
	TUN3	112.187	89	1.261		
	TUN4	118.943	89	1.336		
	TUN5	116.584	89	1.310		
	MMB3	179.926	89	2.022		
	MMB4	178.162	89	2.002		
	MMB6	191.731	89	2.154		
	MMB7	180.887	89	2.032		

Table L.17 – Adjusted Estimates – Learning Modes

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
MMM1	A	4.670(a)	.210	4.252	5.088
	G	4.948(a)	.219	4.513	5.383
	V	4.757(a)	.207	4.347	5.168
MMM3	A	4.289(a)	.211	3.871	4.708
	G	4.485(a)	.219	4.048	4.921
	V	4.412(a)	.207	4.000	4.824
MMM4	A	4.899(a)	.149	4.602	5.196
	G	5.214(a)	.156	4.905	5.523
	V	5.162(a)	.147	4.870	5.454
MMM5	A	4.561(a)	.185	4.194	4.928
	G	5.072(a)	.192	4.690	5.454
	V	4.881(a)	.181	4.520	5.241

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
TUN2	A	4.821(a)	.209	4.406	5.236
	G	4.962(a)	.217	4.530	5.393
	V	4.574(a)	.205	4.166	4.982
TUN3	A	4.760(a)	.198	4.367	5.153
	G	5.231(a)	.206	4.822	5.640
	V	5.050(a)	.194	4.664	5.437
TUN4	A	4.703(a)	.204	4.298	5.107
	G	5.479(a)	.212	5.057	5.900
	V	5.171(a)	.200	4.773	5.569
TUN5	A	4.227(a)	.202	3.826	4.627
	G	4.770(a)	.210	4.353	5.187
	V	4.590(a)	.198	4.196	4.984
MMB3	A	3.540(a)	.250	3.042	4.037
	G	3.659(a)	.261	3.141	4.177
	V	3.365(a)	.246	2.875	3.854
MMB4	A	3.850(a)	.249	3.355	4.345
	G	4.339(a)	.259	3.824	4.854
	V	4.413(a)	.245	3.926	4.900
MMB6	A	3.635(a)	.259	3.121	4.148
	G	3.921(a)	.269	3.386	4.455
	V	3.615(a)	.254	3.110	4.120
MMB7	A	4.214(a)	.251	3.715	4.713
	G	4.178(a)	.261	3.658	4.697
	V	4.184(a)	.247	3.693	4.675

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.18 – Pairwise Comparisons – Learning Modes

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
MMM1	A	G	-.278	.304	1.000	-1.021	.464
		V	-.088	.298	1.000	-.814	.638
	G	A	.278	.304	1.000	-.464	1.021
		V	.191	.306	1.000	-.556	.937
	V	A	.088	.298	1.000	-.638	.814
		G	-.191	.306	1.000	-.937	.556
MMM3	A	G	-.195	.305	1.000	-.940	.550
		V	-.123	.298	1.000	-.851	.605
	G	A	.195	.305	1.000	-.550	.940
		V	.072	.307	1.000	-.676	.821
	V	A	.123	.298	1.000	-.605	.851
		G	-.072	.307	1.000	-.821	.676
MMM4	A	G	-.315	.216	.446	-.843	.213

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
MMM5	G	V	-.263	.211	.649	-.779	.253
		A	.315	.216	.446	-.213	.843
		V	.052	.218	1.000	-.479	.583
	V	A	.263	.211	.649	-.253	.779
		G	-.052	.218	1.000	-.583	.479
		A	-.511	.267	.177	-1.163	.141
TUN2	A	V	-.320	.261	.673	-.957	.318
		A	.511	.267	.177	-.141	1.163
		V	.191	.269	1.000	-.464	.847
	V	A	.320	.261	.673	-.318	.957
		G	-.191	.269	1.000	-.847	.464
		A	-.141	.302	1.000	-.878	.596
TUN3	A	V	.247	.295	1.000	-.474	.967
		A	.141	.302	1.000	-.596	.878
		V	.387	.304	.617	-.354	1.129
	V	A	-.247	.295	1.000	-.967	.474
		G	-.387	.304	.617	-1.129	.354
		A	-.471	.286	.311	-1.169	.228
TUN4	A	V	-.290	.280	.909	-.973	.393
		A	.471	.286	.311	-.228	1.169
		V	.181	.288	1.000	-.521	.883
	G	A	.290	.280	.909	-.393	.973
		G	-.181	.288	1.000	-.883	.521
		V	-.776(*)	.295	.030	-1.495	-.057
TUN5	G	V	-.468	.288	.324	-1.171	.235
		A	.776(*)	.295	.030	.057	1.495
		V	.308	.296	.905	-.415	1.031
	V	A	.468	.288	.324	-.235	1.171
		G	-.308	.296	.905	-1.031	.415
		A	-.543	.292	.198	-1.255	.169
MMB3	A	V	-.363	.285	.620	-1.059	.333
		G	.543	.292	.198	-.169	1.255
		V	.180	.293	1.000	-.536	.896
	G	A	.363	.285	.620	-.333	1.059
		G	-.180	.293	1.000	-.896	.536
		A	-.119	.362	1.000	-1.004	.765
MMB4	G	V	.175	.354	1.000	-.690	1.040
		A	.119	.362	1.000	-.765	1.004
		V	.294	.364	1.000	-.595	1.184
	V	A	-.175	.354	1.000	-1.040	.690
		G	-.294	.364	1.000	-1.184	.595
		A	-.489	.361	.536	-1.369	.391
MMB6	G	V	-.563	.353	.343	-1.423	.298
		A	.489	.361	.536	-.391	1.369
		V	-.074	.363	1.000	-.959	.811
	V	A	.563	.353	.343	-.298	1.423
		G	.074	.363	1.000	-.811	.959
		A	-.286	.374	1.000	-1.199	.627

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
MMB7	G	V	.020	.366	1.000	-.873	.912
		A	.286	.374	1.000	-.627	1.199
		V	.306	.376	1.000	-.612	1.224
	V	A	-.020	.366	1.000	-.912	.873
		G	-.306	.376	1.000	-1.224	.612
		A	.036	.363	1.000	-.850	.923
	A	V	.030	.355	1.000	-.837	.897
		A	-.036	.363	1.000	-.923	.850
		V	-.006	.365	1.000	-.898	.885
	V	A	-.030	.355	1.000	-.897	.837
G		.006	.365	1.000	-.885	.898	

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

L.3 MANCOVA – Full Analysis Alternative Composites

The results of the MANCOVA using the alternative composite variables are presented below in Tables L.19 through L.22. Note – due to transformations, lower values represent higher scores.

Table L.19 – Multivariate Tests of Significance – Alternative Full Analysis

Effect	Wilks' Lambda	F	df	p
TMP31	0.891	0.874	11; 79	0.569
TOTPIIT2	0.823	1.544	11; 79	0.133
TOTGCSE3	0.847	1.297	11; 79	0.242
TOTWSE2	0.941	0.454	11; 79	0.925
AGE	0.909	0.718	11; 79	0.718
YEAR	0.895	0.844	11; 79	0.597
GENDER	0.790	1.905	11; 79	0.051
TOOL	0.548	2.524	22; 158	0.001

Table L.20 – Tests of Between-Subjects Effects – Alternative Full Analysis

Source	Dependent Variable	SS	df	MS	F	p
Intercept	TOTDEEP2	0.134	1	0.134	1.547	0.217
	TOTBROAD3	0.049	1	0.049	0.376	0.541
	TOTNOVEL3	0.328	1	0.328	2.517	0.116
	TOTRELN3	0.318	1	0.318	2.691	0.104
	TOTSTERM	0.835	1	0.835	1.353	0.248
	TOTFORAGE	0.143	1	0.143	0.277	0.600
	TOTINTERP3	1.837	1	1.837	2.979	0.088
	TOTINSIGHT4	0.391	1	0.391	0.631	0.429
	TOTMMM	0.254	1	0.254	0.438	0.510
	TOTTUN	0.284	1	0.284	0.440	0.509
TOTMMB	0.004	1	0.004	0.007	0.935	
TMP31	TOTDEEP2	0.255	1	0.255	2.934	0.090
	TOTBROAD3	0.011	1	0.011	0.084	0.773
	TOTNOVEL3	0.075	1	0.075	0.573	0.451
	TOTRELN3	0.075	1	0.075	0.635	0.428
	TOTSTERM	1.541	1	1.541	2.498	0.118
	TOTFORAGE	1.317	1	1.317	2.539	0.115
	TOTINTERP3	0.180	1	0.180	0.292	0.590
	TOTINSIGHT4	0.284	1	0.284	0.459	0.500
	TOTMMM	1.907	1	1.907	3.289	0.073
	TOTTUN	0.015	1	0.015	0.024	0.878
TOTMMB	0.242	1	0.242	0.396	0.531	
TOTPIIT2	TOTDEEP2	0.303	1	0.303	3.486	0.065
	TOTBROAD3	0.100	1	0.100	0.763	0.385
	TOTNOVEL3	0.000	1	0.000	0.000	0.984
	TOTRELN3	0.035	1	0.035	0.294	0.589
	TOTSTERM	0.107	1	0.107	0.174	0.678
	TOTFORAGE	1.101	1	1.101	2.123	0.149
	TOTINTERP3	1.360	1	1.360	2.205	0.141
	TOTINSIGHT4	1.727	1	1.727	2.789	0.098
	TOTMMM	1.836	1	1.836	3.167	0.079
	TOTTUN	1.430	1	1.430	2.216	0.140
TOTMMB	1.756	1	1.756	2.877	0.093	
TOTGCSE3	TOTDEEP2	0.144	1	0.144	1.663	0.201
	TOTBROAD3	0.003	1	0.003	0.020	0.889
	TOTNOVEL3	0.017	1	0.017	0.131	0.719
	TOTRELN3	0.145	1	0.145	1.225	0.271
	TOTSTERM	0.693	1	0.693	1.123	0.292
	TOTFORAGE	0.094	1	0.094	0.182	0.671
	TOTINTERP3	0.605	1	0.605	0.982	0.324
	TOTINSIGHT4	1.258	1	1.258	2.030	0.158
	TOTMMM	0.404	1	0.404	0.697	0.406
	TOTTUN	0.757	1	0.757	1.173	0.282
TOTMMB	0.073	1	0.073	0.119	0.731	
TOTWSE2	TOTDEEP2	0.011	1	0.011	0.130	0.719

Source	Dependent Variable	SS	df	MS	F	p
	TOTBROAD3	0.039	1	0.039	0.300	0.585
	TOTNOVEL3	0.082	1	0.082	0.630	0.430
	TOTRELN3	0.068	1	0.068	0.578	0.449
	TOTSTERM	0.510	1	0.510	0.826	0.366
	TOTFORAGE	0.008	1	0.008	0.015	0.903
	TOTINTERP3	0.006	1	0.006	0.009	0.923
	TOTINSIGHT4	0.034	1	0.034	0.055	0.815
	TOTMMM	0.100	1	0.100	0.173	0.679
	TOTTUN	0.008	1	0.008	0.012	0.914
	TOTMMB	0.004	1	0.004	0.006	0.938
AGE	TOTDEEP2	0.037	1	0.037	0.429	0.514
	TOTBROAD3	0.276	1	0.276	2.107	0.150
	TOTNOVEL3	0.000	1	0.000	0.000	1.000
	TOTRELN3	0.009	1	0.009	0.078	0.781
	TOTSTERM	0.224	1	0.224	0.363	0.548
	TOTFORAGE	0.253	1	0.253	0.488	0.487
	TOTINTERP3	0.935	1	0.935	1.516	0.221
	TOTINSIGHT4	0.001	1	0.001	0.001	0.970
	TOTMMM	0.142	1	0.142	0.245	0.622
	TOTTUN	0.340	1	0.340	0.526	0.470
	TOTMMB	0.082	1	0.082	0.134	0.715
YEAR	TOTDEEP2	0.000	1	0.000	0.003	0.958
	TOTBROAD3	0.002	1	0.002	0.014	0.906
	TOTNOVEL3	0.154	1	0.154	1.184	0.280
	TOTRELN3	0.161	1	0.161	1.365	0.246
	TOTSTERM	0.314	1	0.314	0.509	0.477
	TOTFORAGE	0.235	1	0.235	0.453	0.503
	TOTINTERP3	2.517	1	2.517	4.081	0.046
	TOTINSIGHT4	0.748	1	0.748	1.208	0.275
	TOTMMM	0.005	1	0.005	0.008	0.929
	TOTTUN	0.090	1	0.090	0.140	0.709
	TOTMMB	0.012	1	0.012	0.020	0.889
GENDER	TOTDEEP2	0.102	1	0.102	1.171	0.282
	TOTBROAD3	1.146	1	1.146	8.753	0.004
	TOTNOVEL3	0.471	1	0.471	3.618	0.060
	TOTRELN3	0.277	1	0.277	2.348	0.129
	TOTSTERM	0.619	1	0.619	1.004	0.319
	TOTFORAGE	3.120	1	3.120	6.014	0.016
	TOTINTERP3	4.250	1	4.250	6.892	0.010
	TOTINSIGHT4	7.523	1	7.523	12.145	0.001
	TOTMMM	0.646	1	0.646	1.114	0.294
	TOTTUN	2.661	1	2.661	4.122	0.045
	TOTMMB	1.739	1	1.739	2.849	0.095
TOOL	TOTDEEP2	0.647	2	0.324	3.727	0.028
	TOTBROAD3	0.768	2	0.384	2.935	0.058
	TOTNOVEL3	0.618	2	0.309	2.370	0.099
	TOTRELN3	0.017	2	0.009	0.073	0.930

Source	Dependent Variable	SS	df	MS	F	p
	TOTSTERM	1.666	2	0.833	1.350	0.265
	TOTFORAGE	19.733	2	9.866	19.021	0.000
	TOTINTERP3	7.713	2	3.856	6.253	0.003
	TOTINSIGHT4	1.267	2	0.633	1.023	0.364
	TOTMMM	1.341	2	0.670	1.156	0.319
	TOTTUN	2.321	2	1.161	1.798	0.172
	TOTMMB	0.250	2	0.125	0.205	0.815
Error	TOTDEEP2	7.727	89	0.087		
	TOTBROAD3	11.651	89	0.131		
	TOTNOVEL3	11.596	89	0.130		
	TOTRELN3	10.512	89	0.118		
	TOTSTERM	54.912	89	0.617		
	TOTFORAGE	46.166	89	0.519		
	TOTINTERP3	54.884	89	0.617		
	TOTINSIGHT4	55.127	89	0.619		
	TOTMMM	51.603	89	0.580		
	TOTTUN	57.458	89	0.646		
	TOTMMB	54.336	89	0.611		

Table L.21 – Adjusted Estimates – Alternative Full Analysis

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
TOTDEEP2	A	1.836(a)	.052	1.733	1.939
	G	1.660(a)	.054	1.553	1.768
	V	1.843(a)	.051	1.742	1.945
TOTBROAD3	A	1.806(a)	.064	1.679	1.933
	G	1.589(a)	.066	1.457	1.721
	V	1.658(a)	.063	1.534	1.783
TOTNOVEL3	A	1.818(a)	.064	1.691	1.944
	G	1.634(a)	.066	1.502	1.765
	V	1.798(a)	.063	1.674	1.923
TOTRELN3	A	1.816(a)	.061	1.696	1.936
	G	1.795(a)	.063	1.670	1.920
	V	1.784(a)	.060	1.666	1.902
TOTSTERM	A	.177(a)	.138	-.098	.452
	G	-.142(a)	.144	-.428	.144
	V	-.041(a)	.136	-.312	.229
TOTFORAGE	A	.524(a)	.127	.272	.776
	G	-.607(a)	.132	-.869	-.344
	V	.043(a)	.125	-.204	.291
TOTINTERP3	A	.399(a)	.138	.124	.674
	G	-.242(a)	.144	-.528	.044
	V	-.161(a)	.136	-.432	.109
TOTINSIGHT4	A	.121(a)	.139	-.155	.396

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
TOTMMM	G	-.162(a)	.144	-.449	.125
	V	.029(a)	.136	-.242	.300
	A	.154(a)	.134	-.113	.421
TOTTUN	G	-.137(a)	.140	-.415	.140
	V	-.023(a)	.132	-.286	.239
	A	.181(a)	.142	-.100	.462
TOTMMB	G	-.207(a)	.147	-.500	.086
	V	.012(a)	.139	-.264	.289
	A	.054(a)	.138	-.219	.328
	G	-.072(a)	.143	-.356	.213
	V	.012(a)	.135	-.257	.281

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.22 – Pairwise Comparisons – Alternative Full Analysis

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
TOTDEEP2	A	G	.176	.075	.064	-.007	.359
		V	-.007	.073	1.000	-.186	.172
	G	A	-.176	.075	.064	-.359	.007
		V	-.183	.076	.052	-.367	.001
	V	A	.007	.073	1.000	-.172	.186
TOTBROAD3	A	G	.183	.076	.052	-.001	.367
		V	.217	.092	.062	-.008	.442
	G	A	.148	.090	.314	-.072	.368
		V	-.217	.092	.062	-.442	.008
	V	A	-.069	.093	1.000	-.296	.157
		G	-.148	.090	.314	-.368	.072
	TOTNOVEL3	A	G	.069	.093	1.000	-.157
V			.184	.092	.145	-.040	.409
G		A	.019	.090	1.000	-.200	.239
		V	-.184	.092	.145	-.409	.040
TOTRELN3	A	G	-.165	.093	.235	-.391	.061
		V	-.019	.090	1.000	-.239	.200
	G	A	.165	.093	.235	-.061	.391
		V	.021	.088	1.000	-.193	.235
	V	A	.032	.086	1.000	-.177	.241
		G	-.021	.088	1.000	-.235	.193
	TOTSTERM	A	V	.011	.088	1.000	-.204
G			-.032	.086	1.000	-.241	.177
G		A	-.011	.088	1.000	-.226	.204
		V	.319	.200	.343	-.169	.808
	G	A	.218	.196	.803	-.259	.696
		V	-.319	.200	.343	-.808	.169
		V	-.101	.201	1.000	-.592	.390

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
TOTFORAGE	V	A	-.218	.196	.803	-.696	.259
		G	.101	.201	1.000	-.390	.592
	A	G	1.131(*)	.184	.000	.683	1.579
		V	.480(*)	.180	.027	.042	.918
		G	-1.131(*)	.184	.000	-1.579	-.683
TOTINTERP3	V	A	-.650(*)	.185	.002	-1.101	-.200
		G	-.480(*)	.180	.027	-.918	-.042
	A	G	.650(*)	.185	.002	.200	1.101
		V	.641(*)	.200	.006	.153	1.129
		V	.560(*)	.196	.016	.082	1.038
TOTINSIGHT4	G	A	-.641(*)	.200	.006	-1.129	-.153
		V	-.081	.201	1.000	-.572	.410
	V	A	-.560(*)	.196	.016	-1.038	-.082
		G	.081	.201	1.000	-.410	.572
		A	.283	.201	.487	-.207	.772
TOTMMM	G	V	.091	.196	1.000	-.387	.570
		A	-.283	.201	.487	-.772	.207
	V	V	-.191	.202	1.000	-.683	.301
		A	-.091	.196	1.000	-.570	.387
		G	.191	.202	1.000	-.301	.683
TOTTUN	A	G	.292	.194	.410	-.182	.765
		V	.177	.190	1.000	-.286	.641
	G	A	-.292	.194	.410	-.765	.182
		V	-.114	.195	1.000	-.590	.362
		A	-.177	.190	1.000	-.641	.286
TOTMMB	A	G	.114	.195	1.000	-.362	.590
		V	.388	.205	.184	-.112	.888
	G	V	.169	.200	1.000	-.320	.657
		A	-.388	.205	.184	-.888	.112
		V	-.219	.206	.870	-.722	.283
TOTMMB	V	A	-.169	.200	1.000	-.657	.320
		G	.219	.206	.870	-.283	.722
	A	G	.126	.199	1.000	-.360	.612
		V	.042	.195	1.000	-.433	.517
		G	-.126	.199	1.000	-.612	.360
TOTMMB	V	V	-.084	.200	1.000	-.572	.405
		A	-.042	.195	1.000	-.517	.433
	G	.084	.200	1.000	-.405	.572	

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

Using the alternative composites there are minimal changes to the support for theory 2 (see Table L.23), minimal changes to the support for theory 3 (see Table L.24) and no changes to the support for theory 1.

Table L.23 – Support for Theory 2 Research Hypotheses (Alternative Composites)

Knowledge Type			Significant findings?
H1b	Deep	Google > Vivisimo, Athens	Yes at p = 0.1 level
H2b	Broad (new)	Google > Vivisimo, Athens	Partially at p = 0.1 (Google > Athens)
H3b	Novel (new)	Google > Vivisimo, Athens	No
H3b	Relationships (new)	Google > Vivisimo, Athens	No
Learning Mode			
H4b	MMM	Google > Vivisimo, Athens	No
H5b	MTUN	Google > Vivisimo, Athens	No
H6b	MMB	Google > Vivisimo, Athens	No
Learning Process			
H7b	Forage	Google > Vivisimo, Athens	Yes at p = 0.01 level
H8b	Insights (new)	Google > Vivisimo, Athens	No
H9b	Interp (new)	Google > Vivisimo, Athens	Partially at p = 0.01

Table L.24 – Support for Theory 3 Research Hypotheses (Alternative Composites)

H₀ – Google will outperform other tools 5 – 15%				
	Construct	5%	10%	15%
Knowledge Type				
H1a	Deep	accept	accept	accept
H2a	Broad (new)	accept	accept	accept (A)/reject (V)
H3a	Novel (new)	accept	accept	accept
H3a	Relationships (new)	accept	reject	reject
Learning Mode				
H4a	MMM	accept	accept	accept
H5a	MTUN	accept	accept	accept
H6a	MMB	accept	accept	accept
Learning Process				
H7a	Forage	accept	accept	accept
H8a	Insights (new)	accept	accept	accept
H9a	Interp (new)	accept	accept	accept
H10a	Interp (new)	accept	accept	accept

L.4 ANCOVA – MP3 Knowledge

As a secondary analysis, an ANCOVA was run to determine whether there were any significant differences between tool groups in terms of how much participants learned about the MP3 subject matter during the course of the experimental task. Note – due to transformations, lower values represent higher scores.

Table L.25 – Tests of Between-Subjects Effects – MP3 Knowledge

Dependent Variable: TMP32a

Source	SS	df	MS	F	p
Corrected Model	0.536	9	0.060	1.633	0.118
Intercept	0.017	1	0.017	0.453	0.502
TMP31	0.013	1	0.013	0.354	0.554
TOTPIIT2	0.013	1	0.013	0.370	0.545
TOTGCSE3	0.104	1	0.104	2.844	0.095
TOTWSE2	0.001	1	0.001	0.038	0.846
AGE	0.004	1	0.004	0.104	0.748
YEAR	0.056	1	0.056	1.544	0.217
GENDER	0.043	1	0.043	1.169	0.283
TOOL	0.212	2	0.106	2.914	0.059
Error	3.244	89	0.036		

Table L.26 – Adjusted Estimates – MP3 Knowledge

Dependent Variable: TMP32a

TOOL	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
A	.498(a)	.034	.431	.565
G	.384(a)	.035	.315	.454
V	.420(a)	.033	.355	.486

a Covariates appearing in the model are evaluated at the following values: TMP31 = 1.5572, TOTPIIT2 = 4.4690, TOTGCSE3 = 7.2088, TOTWSE2 = .0000, AGE = 19.88, YEAR = 2.4848, GENDER = 1.5455.

Table L.27 – Pairwise Comparisons – MP3 Knowledge

Dependent Variable: TMP32a

(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
					Upper Bound	Lower Bound
A	G	.114	.049	.064	-.005	.233
	V	.078	.048	.314	-.038	.194
G	A	-.114	.049	.064	-.233	.005
	V	-.036	.049	1.000	-.156	.083
V	A	-.078	.048	.314	-.194	.038
	G	.036	.049	1.000	-.083	.156

Based on estimated marginal means

a Adjustment for multiple comparisons: Bonferroni.

APPENDIX M EXPERIMENTAL CONTROLS

M.1 Manipulation Check – Quantitative Analysis

The results of a MANOVA to assess the tool functionality manipulation check are reported below.

Table M.1 – Multivariate Tests of Significance – Manipulation Check

Effect	Wilks' Lambda	F	df	p
TOOL	20.660	3.475	12; 182	0.000

Table M.2 – Tests of Between-Subjects Effects – Manipulation Check

Source	Dependent Variable	SS	df	MS	F	p
Intercept	Zscore(TLEVEL11)	0.003	1	0.003	0.003	0.954
	Zscore(TLEVEL12)	0.000	1	0.000	0.000	0.988
	Zscore(TLEVEL21)	0.003	1	0.003	0.003	0.956
	Zscore(LEVEL22)	0.011	1	0.011	0.012	0.913
	Zscore(LEVEL31)	0.001	1	0.001	0.001	0.977
	Zscore(TLEVEL32)	0.006	1	0.006	0.007	0.932
TOOL	Zscore(TLEVEL11)	4.801	2	2.400	2.472	0.090
	Zscore(TLEVEL12)	14.176	2	7.088	8.117	0.001
	Zscore(TLEVEL21)	7.065	2	3.532	3.729	0.028
	Zscore(LEVEL22)	5.253	2	2.627	2.719	0.071
	Zscore(LEVEL31)	1.142	2	0.571	0.566	0.570
	Zscore(TLEVEL32)	15.897	2	7.948	9.294	0.000
Error	Zscore(TLEVEL11)	93.199	96	0.971		
	Zscore(TLEVEL12)	83.824	96	0.873		
	Zscore(TLEVEL21)	90.935	96	0.947		
	Zscore(LEVEL22)	92.747	96	0.966		
	Zscore(LEVEL31)	96.858	96	1.009		
	Zscore(TLEVEL32)	82.103	96	0.855		

Table M.3 – Estimates – Manipulation Check

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Zscore(TLEVEL11)	A	.265	.172	-.076	.605
	G	-.283	.177	-.634	.068
	V	.001	.167	-.330	.331
Zscore(TLEVEL12)	A	.534	.163	.211	.856
	G	-.305	.168	-.638	.028
	V	-.233	.158	-.546	.081
Zscore(TLEVEL21)	A	-.345	.169	-.681	-.008
	G	.315	.175	-.032	.662
	V	.046	.165	-.280	.373
Zscore(LEVEL22)	A	.107	.171	-.233	.447
	G	-.337	.177	-.687	.014
	V	.197	.166	-.132	.527
Zscore(LEVEL31)	A	.127	.175	-.220	.474
	G	.013	.180	-.345	.371
	V	-.131	.170	-.468	.206
Zscore(TLEVEL32)	A	-.520	.161	-.839	-.200
	G	.468	.166	.138	.798
	V	.075	.156	-.235	.386

Table M.4 – Pairwise Comparisons – Manipulation Check

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
Zscore(TLEVEL11)	A	G	.548	.246	.086	-.052	1.149
		V	.264	.239	.816	-.318	.847
	G	A	-.548	.246	.086	-1.149	.052
		V	-.284	.243	.737	-.876	.308
	V	A	-.264	.239	.816	-.847	.318
		G	.284	.243	.737	-.308	.876
Zscore(TLEVEL12)	A	G	.839(*)	.234	.002	.269	1.408
		V	.766(*)	.227	.003	.214	1.319
	G	A	-.839(*)	.234	.002	-1.408	-.269
		V	-.072	.230	1.000	-.634	.489
	V	A	-.766(*)	.227	.003	-1.319	-.214
		G	.072	.230	1.000	-.489	.634
Zscore(TLEVEL21)	A	G	-.659(*)	.243	.024	-1.253	-.066
		V	-.391	.236	.303	-.966	.185
	G	A	.659(*)	.243	.024	.066	1.253
		V	.268	.240	.799	-.316	.853
	V	A	.391	.236	.303	-.185	.966
		G	-.268	.240	.799	-.853	.316
Zscore(LEVEL22)	A	G	.444	.246	.223	-.155	1.043

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
Zscore(LEVEL31)	G	V	-.090	.238	1.000	-.671	.491
		A	-.444	.246	.223	-1.043	.155
		V	-.534	.242	.090	-1.125	.057
	V	A	.090	.238	1.000	-.491	.671
		G	.534	.242	.090	-.057	1.125
		A	.114	.251	1.000	-.498	.726
Zscore(TLEVEL32)	G	V	.258	.244	.875	-.335	.852
		A	-.114	.251	1.000	-.726	.498
		V	.144	.248	1.000	-.459	.748
	V	A	-.258	.244	.875	-.852	.335
		G	-.144	.248	1.000	-.748	.459
		A	-.988(*)	.231	.000	-1.551	-.424
G	V	-.595(*)	.224	.028	-1.142	-.048	
	A	.988(*)	.231	.000	.424	1.551	
	V	.393	.228	.265	-.163	.949	
V	A	.595(*)	.224	.028	.048	1.142	
	G	-.393	.228	.265	-.949	.163	

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

An assessment of the manipulation check results is provided below in Table M.5.

Table M.5 – Manipulation Check – Tool Characteristics

Tool Characteristics	Order of Means	Differences Supported?	Comments
LEVEL11: Searches the Web using keywords entered.	G, V, A.	Partially supported at α 0.1 level (G>A)	Differences between Google and Athens make sense because Athens doesn't do this in a direct way, whereas both Google and Vivisimo do.
LEVEL12: Includes links to results that are most similar to keywords entered.	G, V, A.	Partially supported at α 0.01 (G>A, V>A)	Differences make sense because Athens's links aren't similar to the keywords entered, whereas both Google's and Vivisimo's are.
LEVEL21: Groups results into clusters of similar content.	A, V, G.	Partially supported at α 0.05 (A>G)	Differences make sense and highlight the grouping and clustering capabilities of Athens and Vivisimo. Clustering was emphasized strongly in Athens training and a little bit in Vivisimo's training.
LEVEL22: Allows users to drill down into clusters of similar content.	V, A, G.	Partially supported at α 0.1 (V>G)	Differences make sense. Google doesn't do this, but Vivisimo and Athens do.
LEVEL31: Draws attention to interesting and novel results that are indirectly related to keywords entered.	A, G, V.	No significant differences.	Athens has the highest mean as expected, but the difference is not significant.

Tool Characteristics	Order of Means	Differences Supported?	Comments
LEVEL32: Results are indirectly connected to keywords, rather than most similar to keywords.	A, V, G.	Supported at α 0.001 level (A>G) and p.05 level (A>V)	Differences make sense since this is what Athens does.

M.2 Manipulation Check – Qualitative Analysis

A qualitative assessment of the degree to which participants in each tool group understood the experimental task – level 3 task – are presented below.

M.2.1 Athens Tool Group

In the Athens group, two participants did not respond to this question. Four answers did not contain the word new, novel, innovative, radical, unique or different as it pertained to ideas, but participants still generally seemed to understand the task:

- “In my search, I was trying to find information related to the latest developments in digital music, both software and hardware-related.”
- “In my search, I was attempting to touch on different aspects of everyday life that could be enhanced or improved if combined with the use of MP3s.”
- “Trying to get a good idea, think outside the box, and look for related terms.
- “Was trying to make links between an MP3 player and a potential service which could augment the MP3 experience like wireless headphones.”

The following responses suggest the participants understood the task, but provide some evidence that individuals were searching for something specific and that they had already generated an idea *a priori*:

- “I was attempting to find a novel idea for the application of MP3 technology with a focus on wireless data streaming.”
- “Was trying to make links between an MP3 player and a potential service which could augment the MP3 experience like wireless headphones.”

M.2.2 Google Tool Group

In the Google group, all respondents answered this question. The following answers did not contain the key words, but participants still generally seemed to understand the task:

- “To find an interesting idea that an MP3 company could use to expand their product or service line or market in general.”
- “I was trying to think of technology not currently being produced that would be applicable to a target market similar to my own demographic (18-25, University student, mid-high social class).”
- “Trying to get general information from which I could think of ideas.”
- “Determining the potentials that exist or paths to be taken in the near future market/industry.”
- “I was trying to develop new products using MP3 technology; however, I ended up researching improvements to the already existing products rather than finding new ones.”

For the following responses, I question the degree to which the participants understood the overall task:

- “To find some information about auditory entertainment to help me market an mp3.”
- “Improve product quality as well as the service.”
- “To develop a product that would help the organization.”
- “Through the search, I was trying to find relationships between two ideas and see the connection between them. I wanted links that would have detailed information about my search terms.”

The following response suggests the participant understood the task, but provides some evidence that individuals were searching for something specific and that they had already generated an idea *a priori*. Thus, the tool wasn’t really helping them to generate novel ideas.

- “I was trying to find ideas about novel applications to mp3 players. I was also trying to validate an idea that I already had, which was to create playlists on the go.”

M.2.3 Vivisimo Tool Group

In the Vivisimo group, one participant did not respond to this question. The following answers did not have the key words in them, but participants still generally seemed to understand the task:

- “Relate possible other commercial uses of mp3 technology and players. Other functions that it can be used for. Other kinds of data transmission, etc.”

- “I was trying to find out more information about MP3 players, and the current market potential for MP3 players.”
- “I was trying to get more details about the technology so I can properly apply them to my ideas. At the same time, I was looking for related ideas that can be applied to an MP3 player. I want to make sure that all the features I added satisfy some important needs and to save people's time.”
- “I was attempting to find a company that had recently come out with an amazing idea for MP3s but had not been taken advantage of yet.”
- “I was trying to look for a need consumers have in the market, in which could be implemented on a mp3 player and sold to the public. I was also looking for what companies are trying to do in the upcoming years to analyze the opportunities in the market.”

For the following responses, I question the degree to which the participants understood

the overall task:

- “To have the most relevant information displayed on the first page of the search.”
- “I was looking for a large variety of information regarding a specific topic, the MP3 player.”
- “When searching using Vivisimo, I was trying to find how MP3 players are being used? And what makes one MP3 player stand out from other competitors. I believe understanding these two segments will allow you to develop a successful MP3 player.”

The following responses suggest that the participants understood the task, but provide interesting insights into how different people approach searching.

- “Any new idea that would be clear from what was seen on the results screen. This way I would not have to open each search result and actually read the page in depth. I wanted obvious answers in front of me, which I did not often get, hence why I only had 3 ideas. I am impatient and dislike reading webpages in more detail unless I know they will be 100% useful.”
- “I was trying to find additional information and background facts about ideas that I already had a basic idea about. For 3 or 4 out of 5 of my ideas, I came up with a brief outline of what I was looking for in my head before I began searching. The search tool gave me the initial idea to search for hearing aid possibilities in MP3 players.”
- “I was trying to generate new ideas for the company, but this was hard without specific knowledge of the direction I already wanted to go in, which I did not have.”

M.3 Satisfaction, Training, Instructions

The results of a MANOVA to assess whether there were differences across tool groups in tool satisfaction, training and instructions are reported below.

Table M.6 – Multivariate Tests of Significance – Satisfaction, Training, Instructions

Effect	Wilks' Lambda	F	df	p
TOOL	0.720	2.297	14; 180	0.006

Table M.7 – Tests of Between-Subjects Effects – Satisfaction, Training, Instructions

Source	Dependent Variable	SS	df	MS	F	p
Intercept	Zscore(TSAT2a)	0.005	1	0.005	0.005	0.941
	Zscore(SAT3)	0.005	1	0.005	0.006	0.941
	Zscore(TSAT4a)	0.003	1	0.003	0.003	0.957
	Zscore(TTRAIN1a)	0.000	1	0.000	0.000	1.000
	Zscore(TTRAIN3)	0.001	1	0.001	0.001	0.976
	Zscore(TCLEAR1a)	0.001	1	0.001	0.001	0.974
	Zscore(TCLEAR2a)	0.002	1	0.002	0.002	0.968
TOOL	Zscore(TSAT2a)	14.239	2	7.119	8.160	0.001
	Zscore(SAT3)	14.850	2	7.425	8.573	0.000
	Zscore(TSAT4a)	13.317	2	6.658	7.548	0.001
	Zscore(TTRAIN1a)	3.022	2	1.511	1.527	0.222
	Zscore(TTRAIN3)	0.387	2	0.193	0.190	0.827
	Zscore(TCLEAR1a)	0.740	2	0.370	0.365	0.695
	Zscore(TCLEAR2a)	0.656	2	0.328	0.324	0.724
Error	Zscore(TSAT2a)	83.761	96	0.873		
	Zscore(SAT3)	83.150	96	0.866		
	Zscore(TSAT4a)	84.683	96	0.882		
	Zscore(TTRAIN1a)	94.978	96	0.989		
	Zscore(TTRAIN3)	97.613	96	1.017		
	Zscore(TCLEAR1a)	97.260	96	1.013		
	Zscore(TCLEAR2a)	97.344	96	1.014		

Table M.8 – Estimates – Satisfaction, Training, Instructions

Dependent Variable	TOOL	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Zscore(TSAT2a)*	A	.498	.163	.175	.821
	G	-.432	.168	-.765	-.099
	V	-.087	.158	-.400	.226
Zscore(SAT3)**	A	.510	.162	.188	.832
	G	-.439	.167	-.771	-.107
	V	-.092	.157	-.405	.220
Zscore(TSAT4a)*	A	.498	.163	.173	.822
	G	-.383	.169	-.718	-.048
	V	-.130	.159	-.445	.185
Zscore(TTRAIN1a)*	A	.247	.173	-.097	.591
	G	-.122	.179	-.477	.233
	V	-.125	.168	-.459	.209
Zscore(TTRAIN3)**	A	-.014	.176	-.363	.334
	G	.087	.181	-.272	.447
	V	-.064	.170	-.402	.274
Zscore(TCLEAR1a)*	A	.077	.175	-.271	.424
	G	.049	.181	-.309	.408
	V	-.116	.170	-.454	.222
Zscore(TCLEAR2a)*	A	.006	.175	-.342	.354
	G	.103	.181	-.256	.462
	V	-.097	.170	-.434	.241

* Due to transformations (reflection), lower values represent higher scores.

**Negatively worded.

Table M.9 – Pairwise Comparisons – Satisfaction, Training, Instructions

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Upper Bound	Lower Bound
Zscore(TSAT2a)	A	G	.930(*)	.234	.000	.361	1.499
		V	.585(*)	.227	.034	.033	1.137
	G	A	-.930(*)	.234	.000	-1.499	-.361
		V	-.345	.230	.413	-.906	.216
	V	A	-.585(*)	.227	.034	-1.137	-.033
		G	.345	.230	.413	-.216	.906
Zscore(SAT3)	A	G	.949(*)	.233	.000	.382	1.516
		V	.602(*)	.226	.027	.052	1.153
	G	A	-.949(*)	.233	.000	-1.516	-.382
		V	-.346	.230	.403	-.906	.213
	V	A	-.602(*)	.226	.027	-1.153	-.052
		G	.346	.230	.403	-.213	.906
Zscore(TSAT4a)	A	G	.881(*)	.235	.001	.308	1.453
		V	.628(*)	.228	.021	.072	1.183
	G	-.881(*)	.235	.001	-1.453	-.308	

Dependent Variable	(I) TOOL	(J) TOOL	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
Zscore(TTRAIN1a)		V	-.253	.232	.831	-.818	.311
	V	A	-.628(*)	.228	.021	-1.183	-.072
		G	.253	.232	.831	-.311	.818
	A	G	.369	.249	.424	-.237	.975
		V	.372	.241	.379	-.216	.960
	G	A	-.369	.249	.424	-.975	.237
Zscore(TTRAIN3)		V	.003	.245	1.000	-.595	.601
	V	A	-.372	.241	.379	-.960	.216
		G	-.003	.245	1.000	-.601	.595
	A	G	-.102	.252	1.000	-.716	.513
		V	.050	.245	1.000	-.546	.646
	G	A	.102	.252	1.000	-.513	.716
Zscore(TCLEAR1a)		V	.151	.249	1.000	-.455	.757
	V	A	-.050	.245	1.000	-.646	.546
		G	-.151	.249	1.000	-.757	.455
	A	G	.027	.252	1.000	-.586	.641
		V	.193	.244	1.000	-.403	.788
	G	A	-.027	.252	1.000	-.641	.586
Zscore(TCLEAR2a)		V	.165	.248	1.000	-.440	.770
	V	A	-.193	.244	1.000	-.788	.403
		G	-.165	.248	1.000	-.770	.440
	A	G	-.097	.252	1.000	-.711	.517
		V	.102	.244	1.000	-.493	.698
	G	A	.097	.252	1.000	-.517	.711
	V	.200	.248	1.000	-.406	.805	
	V	A	-.102	.244	1.000	-.698	.493
	V	G	-.200	.248	1.000	-.805	.406

Based on estimated marginal means

* The mean difference is significant at the .05 level.

a Adjustment for multiple comparisons: Bonferroni.

A summary of the differences in satisfaction between Athens and the other tool groups is presented in Table M.10

Table M.10 – Differences in Satisfaction Between Tool Groups

Question	Order of Means	Significant Differences
[SAT2] I found the tool useful for completing today's task.	G, V, A	Significant differences between Athens and Google (α 0.001) and between Athens and Vivisimo (α 0.05)
[SAT3] I found using the tool frustrating.	A, V, G	Significant differences between Athens and Google (α 0.001) and between Athens and Vivisimo (α 0.05)
[SAT4] I was better off using the tool to complete the task than without the tool.	G, V, A	Significant differences between Athens and Google (α 0.001) and between Athens and Vivisimo (α 0.05)

APPENDIX N QUALITATIVE DATA

Tables N.1 through N.3 present a summary of the most prevalent common themes found across tool groups for the tool features, challenges and suggested improvements questions.

Table N.1 – Features

Tool	Sample Comments	Theme (# of instances)
Athens	“The Clusters effectively grouped the search results. I liked the ability to search deeper within a particular cluster.”	Clusters (18)
	“Clustering certain areas of information. For example, there was a cluster on gaming products which helped me find a good idea for HD technology in game products.”	
	“The ability of the tool to group similar topics and keywords together was very helpful and reduced the time wasted by going through each link individually.”	
	“I liked the clusters and how they grouped certain related keywords together. It allowed me to navigate into one, and narrow down my searches with further clusters.”	
	“The clusters, which grouped relevant and similar information together and allowed me to drill down to find even more specific web links.”	
	“The clusters assisted me with today’s task. The clusters made it easy to see where similar words were grouped and allowed me to quickly narrow down which cluster was more interesting.”	
	“The tool was extremely efficient at helping me find sites relevant to my searches, then further helping me drill down into clusters of them to find more specific and beneficial information.”	
	“The fact that the results were grouped into clusters of similar results was very helpful.”	
Vivisimo	“The grouping of topics. It subdivided ideas for me and allowed me to see the different branches of the mp3 market.”	Clusters (18)
	“I liked the idea of clustering your search results into main headings. This allows me to locate my desired information faster and more effectively.”	
	“The grouping features were sometimes helpful to define various areas of results.”	
	“The grouping tool saved me a lot of time and I had a general idea what the results were.”	
Vivisimo	“The advanced web search. I included a lot of the news sites because that’s where new technologies would be featured.”	Advanced search (4)
Google	“The advanced search form in Google helped me to get better results when I was looking for specific information.”	Advanced search (5)
Vivisimo	“Search capability based on keywords I entered.”	Keyword (basic) search (10)
Google	“I mainly used common keyword searches - nothing too complicated. I found that coming up with the right words and phrases to search was the key here.”	Keyword (basic) search (6)

Tool	Sample Comments	Theme (# of instances)
Vivisimo	"The way it presented results based on topics that most related to what I was searching for."	Relevant results (5)
Google	"The tool usually produces relevant results. It also usually tends to produce a Wikipedia link, which was really useful for today's task."	Relevant results (5)

Other interesting results:

Some of the Vivisimo comments included references and comparisons to Google. For example:

"The layout with the category of search results was helpful in completing today's task, it was familiar like Google's layout."

Table N.2 – Challenges

Tool	Sample Comments	Theme (# of instances)
Athens	"This tool does not clearly attack the queries asked of it. Although it provides a broad spectrum of options, much of the time I found myself sifting through unnecessary websites that prolonged my search."	Irrelevant results (13)
	"Some of the webpages were very irrelevant to the topic and had nothing to do with audio technology at all."	
	"Sometimes it went completely off track and brought back results which at least on the face of it seemed completely unrelated topics which had nothing to do with the search."	
Google	"It would often retrieve information/websites that were not relevant to my task. Sometimes it would display results that contained only one of the search terms instead of all of the ones I had used."	Irrelevant results (8)
	"Some websites that the results generated were completely off, for example when I was searching exactly 'digital rights management mp3-players'."	
Athens	"I found it almost no better than a regular search tool, such as Google."	Comparison to Google and other tools (3)
	"Being new to Athens, I required a few minutes to play around on it to better familiarize myself with its setup. Also, I was a little overwhelmed with all of the information at the beginning, while I simply found it different from typical search engines."	
Vivisimo	"I didn't like the fact that they didn't provide me with spelling corrections, unlike Google if I mistyped a word they just searched for it as was and gave me useless results. I also didn't like that you had to already know in what direction you wanted to search before starting."	Comparison to Google and other tools (4)
	"I'm used to Google so my natural inclination is to type search parameters in the Google toolbox in the top right, had to remind myself to use Vivisimo."	
Athens	"Did not give descriptions of website links like traditional search websites (aka - Google) -- resulted in clicking on links that were relevant to the topic."	Lack of metadata (3)
Google	"The summaries were sometimes misleading. Once, the summary gave a relevant idea, yet when I clicked the link, I	Lack of metadata (1)

Tool	Sample Comments	Theme (# of instances)
	was directed to a list of articles. I did not know which article related to the summary I was interested in."	
Athens	"The searches were not what I expected. Because my search words were less-focused, much of results were useless."	Search term definition and refinement (4)
Vivisimo	"I also didn't like that you had to already know in what direction you wanted to search before starting."	Search term definition and refinement (3)
Google	"I had to continue to refine my search terms, as it is not always easy to figure out which words will produce the best results."	Search term definition and refinement (6)
Athens	"The results were irrelevant to what I was looking for."	Didn't give me what I wanted/ was looking for (5)
Vivisimo	"The tool took the keyword but not what I was looking for always. It usually gave me sites that offered Mp3's for sale. I never indicated purchase at all in the search. I wanted information."	Didn't give me what I wanted/ was looking for (13)
Google	"Information was often too general in the webpages and I was unable to find what i was specifically looking for." "I found it harder to find interesting ideas by searching general keywords or phrases."	Didn't give me what I wanted/ was looking for (5)

Table N.3 – Suggested Improvements

Tool	Sample Comments	Theme (# of instances)
Athens	"It is difficult to say, because the tool uses a different approach. I would appreciate, if it provided some of the Google-type search too, because it would be more helpful with the search words I chose."	Comparison to Google (2)
Vivisimo	"Sponsored links provided at the side (similar to Google)."	Comparison to Google (1)
Athens	"Option to filter sources. Be able to choose the level of how abstract / indirect the results are." "Sorting the data (such as by relevance, date, author, etc)." "To be able to refine your search results to articles, journals or what you actually need. I found many of the links were ads for products, which was not overly helpful."	Filtering and sorting (4)
Vivisimo	"Allow user to sort results." "It would be good if there was a way to specify which kind of sites you were looking for - i.e. academic, or sales, or organization."	Filtering and sorting (3)
Google	"In order to improve this tool, Google should allow people to decide whether they want results that either explicitly apply or loosely apply to the search words entered in order to decrease the number of unrelated pages that appear."	Filtering and sorting (3)
Athens	"Including descriptions of website links so that you don't have to click on them to understand their content."	Metadata (3)
Vivisimo	"Summary of what is in each site (which comes below the link)."	Metadata (1)
Google	"Potentially Google could provide more information on the hit (i.e. it comes from a news cite, Magazine article, etc.)"	Metadata (3)

Tool	Sample Comments	Theme (# of instances)
	along with it's 2 line summary of the site that it shows under the web address link."	
Athens	"I think it would be more useful to have the information better organized and categorized in descriptive English (as opposed to terms that one who hasn't been trained might not understand). The word 'cluster' was used for more than one level, and was already somewhat confusing." "Simpler navigation. I found it challenging to find my way around and to relocate results."	Navigation (16)
Vivisimo	"I didn't really like the layout of the tool. I would also allow users to do an advanced search without having to first type in a search. For example, I don't think I could do an advanced search for the first search terms I tried."	Navigation (3)
Google	"More intriguing interface."	Navigation (1)
Athens	"More statistics would be better (e.g. relevancy, number of results / uniqueness)."	Quality indicators (2)
Vivisimo	"Ranking of page would be helpful."	Quality indicators (1)
Google	"Maybe it could show the percentage correlation per link? But I don't really have any issues with the tool."	Quality indicators (2)
Athens	"It should be easier to refine your search. Most people don't have a few hours to wait to search for ideas."	Search term refinement (6)
Vivisimo	"Possibly a thesaurus type program where if you type in a word, it gives you other similar words to try. For example, when I type in "portable music + uses" if a list came up somewhere that showed other possible synonyms like "mp3 player, walkman" for the first word, and "application" for the second word."	Search term refinement (3)
Google	"The most beneficial would be search within a search option. So let's say I got a result I liked, but the other results were poor. You should be able to click something like "similar to this result" which would help you target your search more fluently and efficiently."	Search term refinement (3)

Other interesting results:

Google participants suggested improvements to the relevancy of results and the ability of the tool to understand the context of the search.

APPENDIX O IDEA EVALUATION

The following table provides a summary of the radical and benefit item questions and anchors used by the judges when evaluating the ideas from the lab experiment.

Table O.1 – Summary of Benefit and Radical Items

Item	Question Wording	Anchors (7-point scale)
Rad1	How novel will this idea be to the organization, based on what the organization does now?	1 - Slightly new 4 - Moderately new 7 - Radically new
Rad2	What will the impact of this idea be on the organization's products or services?	1 - Improve existing product functionality or services 4 - Add new functionality or dimensions to existing products/services 7 - Create a new product or service
Rad4	What will the impact of this idea be on the organization's technology?	1 - Use the company's existing technology 4 - Modify the company's existing technology 7 - Use a new technology in the company
Ben1	How profitable will this idea likely be?	1 - Not profitable 4 - Marginally profitable 7 - Highly profitable
Ben3	How will this idea impact the organization's competitiveness?	1 - Not help the company achieve competitive success 4 - Help company achieve short-term competitive success 7 - Help company achieve long-term competitive success
Ben5	How likely is it that the organization will pursue this idea?	1 - Not at all likely 4 - Moderately likely 7 - Highly likely

O.1 Preliminary Analysis

The results of the preliminary analyses performed on the idea evaluation data are discussed below.

O.1.1 Missing Values Analysis:

A missing value analysis was performed in SPSS by question type and by rater (judge) (see Tables O.2 and O.3).

Table O.2 – Missing Values and N/A Entries by Question

Question	Missing	N/A
Ben1	7	38
Ben3	2	39
Ben5	0	42
Rad1	2	42
Rad2	6	47
Rad4	0	38
Total	17	246

Table O.3 – Missing Values and N/A Entries by Rater

Rater	1	2	3	4	5	Total
N/A	184	21	7	0	34	246
Missing	0	8	9	0	0	17

Compared to the other raters, rater 1 had a significant number of N/A entries (184). Thus, this rater’s scores may be less reliable than the other raters.

O.1.2 Reliability

To assess the level of agreement across raters, intraclass correlations were run for each question type. I assessed the intraclass correlation, inter-item correlations (correlations between raters) and Cronbach’s alpha. Overall, there was very low consistency in scores across raters. Consistency was lower for the benefit questions than the radical questions (see summary in Table O.4).

Table O.4 – Reliability Summary by Question Type

Item	ICC (single measures)	Cronbach’s alpha
Ben1	0.1073	0.375
Ben3	0.0535	0.220
Ben5	0.1182	0.401
Rad1	0.2033	0.561
Rad2	0.1943	0.547
Rad4	0.3437	0.724

O.1.2.1 Reliability Assessment – Question Ben1

Table O.5 – Intraclass Correlation Coefficient (Ben1)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0			
		Lower Bound	Upper Bound	Value	df1	df2	Sig
Single Measures	0.1073 (b)	0.0225	0.2192	1.6011	66	264	0.005
Average Measures	0.3754 (c)	0.1032	0.5840	1.6011	66	264	0.005

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.6 – Inter-Item Correlation Matrix (Ben1)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.119	.114	-.014	.122
rater2	.119	1.000	.147	.282	.210
rater3	.114	.147	1.000	.008	.016
rater4	-.014	.282	.008	1.000	.186
rater5	.122	.210	.016	.186	1.000

Table O.7 – Reliability Statistics (Ben1)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.375	.403	5

O.1.2.2 Reliability Assessment – Question Ben3

Table O.8 – Intraclass Correlation Coefficient (Ben3)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0	df1	df2	Sig
		Lower Bound	Upper Bound	Value			
Single Measures	0.0535 (b)	-0.0199	0.1528	1.2829	70	280	0.083
Average Measures	0.2205 (c)	-0.1081	0.4742	1.2829	70	280	0.083

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.9 – Inter-Item Correlation Matrix (Ben3)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.119	.117	-.046	.084
rater2	.119	1.000	.266	.114	-.005
rater3	.117	.266	1.000	-.001	-.192
rater4	-.046	.114	-.001	1.000	.064
rater5	.084	-.005	-.192	.064	1.000

Table O.10 – Reliability Statistics (Ben3)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.220	.215	5

O.1.2.3 Reliability Assessment – Question Ben5

Table O.11 – Intraclass Correlation Coefficient (Ben5)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0	df1	df2	Sig
		Lower Bound	Upper Bound	Value			
Single Measures	0.1182 (b)	0.0328	0.2298	1.6705	68	272	0.002
Average Measures	0.4014 (c)	0.1449	0.5987	1.6705	68	272	0.002

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.12 – Inter-Item Correlation Matrix (Ben5)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.115	.031	-.007	.135
rater2	.115	1.000	.122	.200	.189
rater3	.031	.122	1.000	.203	.089
rater4	-.007	.200	.203	1.000	.142
rater5	.135	.189	.089	.142	1.000

Table O.13 – Reliability Statistics (Ben5)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.401	.410	5

O.1.2.4 Reliability Assessment – Question Rad1

Table O.14 – Intraclass Correlation Coefficient (Rad1)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0	df1	df2	Sig
		Lower Bound	Upper Bound	Value			
Single Measures	0.2033 (b)	0.1067	0.3229	2.2760	69	276	0.000
Average Measures	0.5606 (c)	0.3739	0.7046	2.2760	69	276	0.000

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.15 – Inter-Item Correlation Matrix (Rad1)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.207	-.055	.003	.053
rater2	.207	1.000	.378	.449	.298
rater3	-.055	.378	1.000	.324	.217
rater4	.003	.449	.324	1.000	.336
rater5	.053	.298	.217	.336	1.000

Table O.16 – Reliability Statistics (Rad1)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.561	.586	5

O.1.2.5 Reliability Assessment – Question Rad2

Table O.17 – Intraclass Correlation Coefficient (Rad2)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0			
		Lower Bound	Upper Bound	Value	df1	df2	Sig
Single Measures	0.1943 (b)	0.0955	0.3183	2.2058	64	256	0.000
Average Measures	0.5466 (c)	0.3456	0.7001	2.2058	64	256	0.000

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.18 – Inter-Item Correlation Matrix (Rad2)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.234	-.124	-.036	.049
rater2	.234	1.000	.299	.387	.347
rater3	-.124	.299	1.000	.254	.147
rater4	-.036	.387	.254	1.000	.462
rater5	.049	.347	.147	.462	1.000

Table O.19 – Reliability Statistics (Rad2)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.547	.559	5

O.1.2.6 Reliability Assessment – Question Rad4

Table O.20 – Intraclass Correlation Coefficient (Rad4)

	Intraclass Correlation(a)	95% Confidence Interval		F Test with True Value 0	df1	df2	Sig
		Lower Bound	Upper Bound	Value			
Single Measures	0.3437 (b)	0.2368	0.4651	3.6180	71	284	0.000
Average Measures	0.7236 (c)	0.6080	0.8130	3.6180	71	284	0.000

Two-way mixed effects model where people effects are random and measures effects are fixed.

a Type C intraclass correlation coefficients using a consistency definition-the between-measure variance is excluded from the denominator variance.

b The estimator is the same, whether the interaction effect is present or not.

c This estimate is computed assuming the interaction effect is absent, because it is not estimable otherwise.

Table O.21 – Inter-Item Correlation Matrix (Rad4)

	rater1	rater2	rater3	rater4	rater5
rater1	1.000	.080	.177	.087	.202
rater2	.080	1.000	.443	.413	.431
rater3	.177	.443	1.000	.559	.559
rater4	.087	.413	.559	1.000	.441
rater5	.202	.431	.559	.441	1.000

Table O.22 – Reliability Statistics (Rad4)

Cronbach's Alpha	Cronbach's Alpha Based on Standardized Items	N of Items
.724	.720	5

O.2 Main Analysis

In preparation for the main analysis, I calculated the mean score for each question, for each idea across the five raters. To assess the degree to which there were differences in mean scores across tools, analysis of variance (ANOVA) was used. Since there is variability in the question type, I grouped the data by question type to see if there were differences between tool groups within a question type. There were no significant differences. However, it is interesting to note that for the radical questions, Athens had consistently higher mean scores (just not significant).

Table O.23 – Tests of Between-Subjects Effects – Idea Scores

Dependent Variable: mean score						
Question	Source	SS	df	MS	F	p
ben1	Corrected Model	0.315	2	0.158	0.201	0.819
	Intercept	1120.485	1	1120.485	1424.905	0.000
	Tool	0.315	2	0.158	0.201	0.819
	Error	80.209	102	0.786		
ben3	Corrected Model	0.843	2	0.422	0.634	0.533
	Intercept	1141.097	1	1141.097	1715.119	0.000
	Tool	0.843	2	0.422	0.634	0.533
	Error	67.862	102	0.665		
ben5	Corrected Model	0.085	2	0.042	0.055	0.946
	Intercept	991.683	1	991.683	1294.324	0.000
	Tool	0.085	2	0.042	0.055	0.946
	Error	78.150	102	0.766		
rad1	Corrected Model	0.715	2	0.358	0.267	0.766
	Intercept	1770.940	1	1770.940	1323.955	0.000
	Tool	0.715	2	0.358	0.267	0.766
	Error	136.437	102	1.338		
rad2	Corrected Model	3.895	2	1.948	1.127	0.328
	Intercept	2028.486	1	2028.486	1173.489	0.000
	Tool	3.895	2	1.948	1.127	0.328
	Error	176.317	102	1.729		
rad4	Corrected Model	3.082	2	1.541	0.687	0.505
	Intercept	1932.564	1	1932.564	861.696	0.000
	Tool	3.082	2	1.541	0.687	0.505
	Error	228.760	102	2.243		

Table O.24 – Estimates – Idea Scores

Dependent Variable: mean score

Question	Tool	Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound	Lower Bound	Upper Bound
ben1	A	3.265	.148	2.972	3.558
	G	3.345	.157	3.034	3.656
	V	3.209	.146	2.920	3.499
ben3	A	3.224	.136	2.954	3.493
	G	3.432	.144	3.146	3.718
	V	3.254	.134	2.988	3.520
ben5	A	3.048	.146	2.758	3.337
	G	3.118	.155	2.811	3.425
	V	3.073	.144	2.787	3.358
rad1	A	4.231	.193	3.849	4.613
	G	4.053	.204	3.648	4.459
	V	4.061	.190	3.684	4.438
rad2	A	4.644	.219	4.210	5.079
	G	4.384	.232	3.923	4.845
	V	4.183	.216	3.755	4.612
rad4	A	4.506	.250	4.010	5.001
	G	4.296	.265	3.771	4.821
	V	4.095	.246	3.606	4.583

O.2.1 ANOVA's by Rater

Given there was low inter-rater reliability, I assessed the degree to which there were any differences between tool groups within a question type at the individual rater level. Thus, I conducted ANOVA's with the rater's scores as the dependent variable, not the mean scores.

O.2.1.1 ANOVA Rater 1

For rater 1, the omnibus test indicates significant differences between tool groups for the question Rad4 (α 0.10). However, in the pairwise comparison, there are no significant differences¹.

¹ Close to significant differences at the p 0.1 level (A > V).

Table O.25 – Tests of Between-Subjects Effects – Idea Scores: Rater 1
Dependent Variable: Rater1

Question	Source	SS	df	MS	F	p
ben1	Corrected Model	0.849	2	0.424	0.132	0.877
	Intercept	973.303	1	973.303	301.948	0.000
	Tool	0.849	2	0.424	0.132	0.877
	Error	235.309	73	3.223		
ben3	Corrected Model	0.334	2	0.167	0.053	0.948
	Intercept	773.200	1	773.200	245.447	0.000
	Tool	0.334	2	0.167	0.053	0.948
	Error	226.812	72	3.150		
ben5	Corrected Model	0.666	2	0.333	0.159	0.854
	Intercept	699.592	1	699.592	333.706	0.000
	Tool	0.666	2	0.333	0.159	0.854
	Error	144.654	69	2.096		
rad1	Corrected Model	0.865	2	0.433	0.124	0.884
	Intercept	917.409	1	917.409	262.324	0.000
	Tool	0.865	2	0.433	0.124	0.884
	Error	251.801	72	3.497		
rad2	Corrected Model	2.152	2	1.076	0.298	0.743
	Intercept	1265.727	1	1265.727	350.255	0.000
	Tool	2.152	2	1.076	0.298	0.743
	Error	249.348	69	3.614		
rad4	Corrected Model	17.967	2	8.983	2.383	0.099
	Intercept	817.737	1	817.737	216.963	0.000
	Tool	17.967	2	8.983	2.383	0.099
	Error	275.139	73	3.769		

Table O.26 – Estimates – Idea Scores: Rater 1
Dependent Variable: Rater1

Question	Tool	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
ben1	A	3.483	.333	2.818	4.147
	G	3.619	.392	2.838	4.400
	V	3.731	.352	3.029	4.433
ben3	A	3.179	.335	2.510	3.847
	G	3.333	.387	2.561	4.105
	V	3.192	.348	2.498	3.886
ben5	A	3.192	.284	2.626	3.759
	G	3.000	.324	2.354	3.646
	V	3.231	.284	2.664	3.797
rad1	A	3.536	.353	2.831	4.240
	G	3.381	.408	2.567	4.194
	V	3.654	.367	2.923	4.385
rad2	A	4.464	.359	3.748	5.181
	G	4.150	.425	3.302	4.998
	V	4.083	.388	3.309	4.857

Question	Tool	Mean	Std. Error	95% Confidence Interval	
rad4	A	3.931	.361	3.213	4.650
	G	3.190	.424	2.346	4.035
	V	2.808	.381	2.049	3.567

Table O.27 – Pairwise Comparisons – Idea Scores: Rater 1
Dependent Variable: Rater1

Question	(I) Tool	(J) Tool	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Lower Bound	Upper Bound
ben1	A	G	-.136	.514	1.000	-1.397	1.124
		V	-.248	.485	1.000	-1.436	.940
	G	A	.136	.514	1.000	-1.124	1.397
		V	-.112	.527	1.000	-1.402	1.179
	V	A	.248	.485	1.000	-.940	1.436
		G	.112	.527	1.000	-1.179	1.402
ben3	A	G	-.155	.512	1.000	-1.411	1.101
		V	-.014	.483	1.000	-1.199	1.171
	G	A	.155	.512	1.000	-1.101	1.411
		V	.141	.521	1.000	-1.135	1.417
	V	A	.014	.483	1.000	-1.171	1.199
		G	-.141	.521	1.000	-1.417	1.135
ben5	A	G	.192	.431	1.000	-.864	1.249
		V	-.038	.402	1.000	-1.024	.947
	G	A	-.192	.431	1.000	-1.249	.864
		V	-.231	.431	1.000	-1.287	.826
	V	A	.038	.402	1.000	-.947	1.024
		G	.231	.431	1.000	-.826	1.287
rad1	A	G	.155	.540	1.000	-1.169	1.478
		V	-.118	.509	1.000	-1.367	1.130
	G	A	-.155	.540	1.000	-1.478	1.169
		V	-.273	.549	1.000	-1.618	1.072
	V	A	.118	.509	1.000	-1.130	1.367
		G	.273	.549	1.000	-1.072	1.618
rad2	A	G	.314	.557	1.000	-1.051	1.680
		V	.381	.529	1.000	-.917	1.679
	G	A	-.314	.557	1.000	-1.680	1.051
		V	.067	.576	1.000	-1.346	1.479
	V	A	-.381	.529	1.000	-1.679	.917
		G	-.067	.576	1.000	-1.479	1.346
rad4	A	G	.741	.556	.562	-.623	2.104
		V	1.123	.524	.106	-.161	2.408
	G	A	-.741	.556	.562	-2.104	.623
		V	.383	.570	1.000	-1.013	1.779
	V	A	-1.123	.524	.106	-2.408	.161
		G	-.383	.570	1.000	-1.779	1.013

Based on estimated marginal means
a Adjustment for multiple comparisons: Bonferroni.

O.2.1.2 ANOVA Rater 2

No significant differences were detected between tool groups for Rater 2.

Table O.28 – Tests of Between-Subjects Effects – Idea Scores: Rater 2
Dependent Variable: Rater2

Question	Source	SS	df	MS	F	p
ben1	Corrected Model	5.326	2	2.663	1.365	0.260
	Intercept	1186.510	1	1186.510	608.139	0.000
	Tool	5.326	2	2.663	1.365	0.260
	Error	187.301	96	1.951		
ben3	Corrected Model	5.516	2	2.758	1.037	0.358
	Intercept	1485.740	1	1485.740	558.671	0.000
	Tool	5.516	2	2.758	1.037	0.358
	Error	260.623	98	2.659		
ben5	Corrected Model	0.789	2	0.395	0.200	0.819
	Intercept	1063.758	1	1063.758	539.773	0.000
	Tool	0.789	2	0.395	0.200	0.819
	Error	197.075	100	1.971		
rad1	Corrected Model	0.231	2	0.116	0.046	0.955
	Intercept	701.877	1	701.877	280.729	0.000
	Tool	0.231	2	0.116	0.046	0.955
	Error	242.519	97	2.500		
rad2	Corrected Model	2.231	2	1.116	0.284	0.754
	Intercept	1909.202	1	1909.202	485.425	0.000
	Tool	2.231	2	1.116	0.284	0.754
	Error	369.707	94	3.933		
rad4	Corrected Model	1.902	2	0.951	0.209	0.811
	Intercept	1729.517	1	1729.517	380.993	0.000
	Tool	1.902	2	0.951	0.209	0.811
	Error	444.871	98	4.539		

Table O.29 – Estimates – Idea Scores: Rater 2
Dependent Variable: Rater2

Question	Tool	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
ben1	A	3.735	.240	3.260	4.211
	G	3.484	.251	2.986	3.982
	V	3.176	.240	2.701	3.652
ben3	A	4.121	.284	3.558	4.685
	G	3.844	.288	3.272	4.416
	V	3.556	.272	3.016	4.095
ben5	A	3.294	.241	2.816	3.772
	G	3.094	.248	2.601	3.586
	V	3.270	.231	2.812	3.728
rad1	A	2.606	.275	2.060	3.152
	G	2.719	.280	2.164	3.274
	V	2.629	.267	2.098	3.159
rad2	A	4.250	.351	3.554	4.946
	G	4.452	.356	3.744	5.159
	V	4.618	.340	3.942	5.293
rad4	A	4.313	.377	3.565	5.060
	G	3.969	.377	3.221	4.716
	V	4.162	.350	3.467	4.857

O.2.1.3 ANOVA Rater 3

For rater 3, the omnibus test indicates significant differences between tool groups for the question Rad2 (α 0.05). In the pairwise comparison, there are significant differences at the α 0.10 level (A>V, G>V).

Table O.30 – Tests of Between-Subjects Effects – Idea Scores: Rater 3
Dependent Variable: Rater3

Question	Source	SS	df	MS	F	p
ben1	Corrected Model	4.016	2	2.008	0.620	0.540
	Intercept	1243.127	1	1243.127	384.030	0.000
	Tool	4.016	2	2.008	0.620	0.540
	Error	317.232	98	3.237		
ben3	Corrected Model	4.379	2	2.190	0.841	0.434
	Intercept	1492.158	1	1492.158	573.106	0.000
	Tool	4.379	2	2.190	0.841	0.434
	Error	262.967	101	2.604		
ben5	Corrected Model	5.128	2	2.564	1.273	0.284
	Intercept	1146.777	1	1146.777	569.441	0.000
	Tool	5.128	2	2.564	1.273	0.284

	Error	203.400	101	2.014		
rad1	Corrected Model	0.997	2	0.499	0.138	0.871
	Intercept	1903.758	1	1903.758	526.945	0.000
	Tool	0.997	2	0.499	0.138	0.871
	Error	357.669	99	3.613		
rad2	Corrected Model	29.436	2	14.718	3.772	0.026
	Intercept	1484.086	1	1484.086	380.348	0.000
	Tool	29.436	2	14.718	3.772	0.026
	Error	374.584	96	3.902		
rad4	Corrected Model	4.162	2	2.081	0.442	0.644
	Intercept	2139.030	1	2139.030	454.171	0.000
	Tool	4.162	2	2.081	0.442	0.644
	Error	475.684	101	4.710		

Table O.31 – Estimates – Idea Scores: Rater 3
Dependent Variable: Rater3

Question	Tool	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
ben1	A	3.444	.300	2.849	4.040
	G	3.800	.328	3.148	4.452
	V	3.314	.304	2.711	3.918
ben3	A	3.528	.269	2.994	4.061
	G	4.032	.290	3.457	4.607
	V	3.838	.265	3.312	4.364
ben5	A	3.417	.237	2.947	3.886
	G	3.548	.255	3.043	4.054
	V	3.027	.233	2.564	3.490
rad1	A	4.200	.321	3.563	4.837
	G	4.433	.347	3.745	5.122
	V	4.378	.312	3.758	4.998
rad2	A	4.229	.334	3.566	4.891
	G	4.321	.373	3.580	5.062
	V	3.139	.329	2.485	3.792
rad4	A	4.667	.362	3.949	5.384
	G	4.710	.390	3.936	5.483
	V	4.270	.357	3.563	4.978

Table O.32 – Pairwise Comparisons – Idea Scores: Rater 3

Dependent Variable: Rater3

Question	(I) Tool	(J) Tool	Mean Difference (I-J)	Std. Error	Sig.(a)	95% Confidence Interval for Difference(a)	
						Lower Bound	Upper Bound
ben1	A	G	-.356	.445	1.000	-1.439	.728
		V	.130	.427	1.000	-.910	1.170
	G	A	.356	.445	1.000	-.728	1.439
		V	.486	.448	.842	-.605	1.576
	V	A	-.130	.427	1.000	-1.170	.910
		G	-.486	.448	.842	-1.576	.605
ben3	A	G	-.504	.395	.615	-1.467	.458
		V	-.310	.378	1.000	-1.230	.610
	G	A	.504	.395	.615	-.458	1.467
		V	.194	.393	1.000	-.762	1.151
	V	A	.310	.378	1.000	-.610	1.230
		G	-.194	.393	1.000	-1.151	.762
ben5	A	G	-.132	.348	1.000	-.978	.715
		V	.390	.332	.731	-.419	1.198
	G	A	.132	.348	1.000	-.715	.978
		V	.521	.346	.403	-.320	1.363
	V	A	-.390	.332	.731	-1.198	.419
		G	-.521	.346	.403	-1.363	.320
rad1	A	G	-.233	.473	1.000	-1.385	.918
		V	-.178	.448	1.000	-1.270	.913
	G	A	.233	.473	1.000	-.918	1.385
		V	.055	.467	1.000	-1.082	1.192
	V	A	.178	.448	1.000	-.913	1.270
		G	-.055	.467	1.000	-1.192	1.082
rad2	A	G	-.093	.501	1.000	-1.313	1.128
		V	1.090	.469	.067	-.053	2.232
	G	A	.093	.501	1.000	-1.128	1.313
		V	1.183	.498	.058	-.030	2.395
	V	A	-1.090	.469	.067	-2.232	.053
		G	-1.183	.498	.058	-2.395	.030
rad4	A	G	-.043	.532	1.000	-1.338	1.252
		V	.396	.508	1.000	-.840	1.633
	G	A	.043	.532	1.000	-1.252	1.338
		V	.439	.528	1.000	-.847	1.726
	V	A	-.396	.508	1.000	-1.633	.840
		G	-.439	.528	1.000	-1.726	.847

Based on estimated marginal means

a Adjustment for multiple comparisons: Bonferroni.

O.2.1.4 ANOVA Rater 4

No significant differences were detected between tool groups for Rater 4.

Table O.33 – Tests of Between-Subjects Effects – Idea Scores: Rater 4
Dependent Variable: Rater4

Question	Source	SS	df	MS	F	p
ben1	Corrected Model	0.498	2	0.249	0.087	0.917
	Intercept	1241.375	1	1241.375	434.429	0.000
	Tool	0.498	2	0.249	0.087	0.917
	Error	291.463	102	2.857		
ben3	Corrected Model	0.727	2	0.363	0.115	0.892
	Intercept	1264.839	1	1264.839	398.850	0.000
	Tool	0.727	2	0.363	0.115	0.892
	Error	323.463	102	3.171		
ben5	Corrected Model	2.095	2	1.048	0.261	0.771
	Intercept	1390.601	1	1390.601	346.067	0.000
	Tool	2.095	2	1.048	0.261	0.771
	Error	409.867	102	4.018		
rad1	Corrected Model	5.867	2	2.933	1.127	0.328
	Intercept	2280.978	1	2280.978	876.418	0.000
	Tool	5.867	2	2.933	1.127	0.328
	Error	265.467	102	2.603		
rad2	Corrected Model	5.832	2	2.916	0.899	0.410
	Intercept	2242.942	1	2242.942	691.842	0.000
	Tool	5.832	2	2.916	0.899	0.410
	Error	330.682	102	3.242		
rad4	Corrected Model	6.421	2	3.210	0.769	0.466
	Intercept	2154.861	1	2154.861	516.232	0.000
	Tool	6.421	2	3.210	0.769	0.466
	Error	425.770	102	4.174		

Table O.34 – Estimates – Idea Scores: Rater 4
Dependent Variable: Rater4

Question	Tool	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
ben1	A	3.389	.282	2.830	3.948
	G	3.406	.299	2.814	3.999
	V	3.541	.278	2.989	4.092
ben3	A	3.361	.297	2.772	3.950
	G	3.531	.315	2.907	4.156
	V	3.541	.293	2.960	4.121
ben5	A	3.472	.334	2.810	4.135
	G	3.656	.354	2.953	4.359
	V	3.811	.330	3.157	4.464

Question	Tool	Mean	Std. Error	95% Confidence Interval	
rad1	A	4.944	.269	4.411	5.478
rad2	G	4.688	.285	4.122	5.253
	V	4.378	.265	3.852	4.904
	A	4.917	.300	4.321	5.512
rad4	G	4.625	.318	3.994	5.256
	V	4.351	.296	3.764	4.938
	A	4.472	.341	3.797	5.148
rad4	G	4.875	.361	4.159	5.591
	V	4.270	.336	3.604	4.936
	A	4.472	.341	3.797	5.148

O.2.1.5 ANOVA Rater 5

No significant differences were detected between tool groups for Rater 5.

Table O.35 – Tests of Between-Subjects Effects – Idea Scores: Rater 5
Dependent Variable: Rater5

Question	Source	SS	df	MS	F	p
ben1	Corrected Model	0.468	2	0.234	0.116	0.891
	Intercept	633.206	1	633.206	313.083	0.000
	Tool	0.468	2	0.234	0.116	0.891
	Error	194.159	96	2.022		
ben3	Corrected Model	0.879	2	0.440	0.261	0.771
	Intercept	484.114	1	484.114	287.475	0.000
	Tool	0.879	2	0.440	0.261	0.771
	Error	161.666	96	1.684		
ben5	Corrected Model	0.526	2	0.263	0.138	0.871
	Intercept	419.346	1	419.346	219.852	0.000
	Tool	0.526	2	0.263	0.138	0.871
	Error	183.110	96	1.907		
rad1	Corrected Model	3.282	2	1.641	0.584	0.560
	Intercept	2810.443	1	2810.443	1000.649	0.000
	Tool	3.282	2	1.641	0.584	0.560
	Error	269.628	96	2.809		
rad2	Corrected Model	1.756	2	0.878	0.233	0.793
	Intercept	2453.803	1	2453.803	650.367	0.000
	Tool	1.756	2	0.878	0.233	0.793
	Error	362.203	96	3.773		
rad4	Corrected Model	0.568	2	0.284	0.059	0.942
	Intercept	2194.760	1	2194.760	458.320	0.000
	Tool	0.568	2	0.284	0.059	0.942
	Error	469.293	98	4.789		

Table O.36 – Estimates – Idea Scores: Rater 5
Dependent Variable: Rater5

Question	Tool	Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound	Lower Bound	Upper Bound
ben1	A	2.444	.237	1.974	2.915
	G	2.586	.264	2.062	3.110
	V	2.588	.244	2.104	3.072
ben3	A	2.111	.216	1.682	2.540
	G	2.345	.241	1.866	2.823
	V	2.206	.223	1.764	2.648
ben5	A	2.028	.230	1.571	2.485
	G	2.172	.256	1.663	2.681
	V	2.000	.237	1.530	2.470
rad1	A	5.389	.279	4.834	5.943
	G	5.103	.311	4.486	5.721
	V	5.559	.287	4.988	6.129
rad2	A	5.194	.324	4.552	5.837
	G	4.893	.367	4.164	5.622
	V	4.943	.328	4.291	5.595
rad4	A	4.750	.365	4.026	5.474
	G	4.724	.406	3.918	5.531
	V	4.583	.365	3.860	5.307

O.2.2 MANOVA

I calculated a composite benefit score, a composite radical score and a total score for each idea. I used three different techniques and compared the results across techniques to ensure the approach to creating a composite score didn't affect the results in a material way. The three approaches are described below in Table O.37.

Table O.37 – Idea Composite Score Creation Techniques

Approach	Description
Average Scores	Created an average benefit score , and an average radical score for each idea by summing the mean scores for each benefit question and dividing by the number of questions. The same approach was taken for the radical questions. The total score was developed by aggregating these two scores (summing them).
Total Scores	Created a total benefit score and a total radical score for each idea by summing all of the benefit scores across raters and benefit questions. The same approach was taken for radical scores. The total score was developed by aggregating these two scores (summing them). The downside of this approach is that ideas are penalized (lower score) and scores biased when there are missing values.
Sum of Mean Scores	Lastly, I created a summed average benefit score and a summed average radical score for each idea by simply summing the mean scores for each benefit question. The same approach was taken for the radical questions. The total score was developed by aggregating these two scores (summing them).

A MANOVA was performed using all nine scores – benefit scores, radical scores and total scores – as the dependent variables. There were no significant differences in scores between the tool groups. However, looking at the order of tools in terms of average group scores provides some interesting insights. The Athens participants were the top performers for radical scores in all approaches to radical scoring. In all total scores, Athens participants were also the highest scorers. Google participants were the highest performers on the benefit scores and were always second to Athens participants on the other scores. Vivisimo participants were always last.

Table O.38 – Multivariate Tests of Significance – Idea Composite Scores

Effect	Wilks' Lambda	F	df	p
TOOL	0.946	0.696	8; 198	0.695

Table O.39 – Tests of Between-Subjects Effects – Idea Composite Scores

Source	Dependent Variable	SS	df	MS	F	p
Intercept	benefit score	1083.393	1	1083.393	1790.022	0.000
	radical score	1909.171	1	1909.171	1439.917	0.000
	total	5868.938	1	5868.938	3059.145	0.000
	benefit sum	208038.018	1	208038.018	1199.455	0.000
	radical sum	362928.087	1	362928.087	1156.212	0.000
	total sum	1120521.708	1	1120521.708	1794.991	0.000
	benefit score sums	9750.533	1	9750.533	1790.022	0.000
	radical score sums	17182.541	1	17182.541	1439.917	0.000
	total score sums	52820.439	1	52820.439	3059.145	0.000
Tool	benefit score	0.317	2	0.158	0.262	0.770
	radical score	2.238	2	1.119	0.844	0.433
	total	2.345	2	1.173	0.611	0.545
	benefit sum	30.120	2	15.060	0.087	0.917
	radical sum	497.342	2	248.671	0.792	0.456
	total sum	744.774	2	372.387	0.597	0.553
	benefit score sums	2.851	2	1.425	0.262	0.770
	radical score sums	20.139	2	10.070	0.844	0.433
	total score sums	21.109	2	10.554	0.611	0.545
Error	benefit score	61.734	102	0.605		
	radical score	135.241	102	1.326		
	total	195.686	102	1.918		
	benefit sum	17691.271	102	173.444		
	radical sum	32017.192	102	313.894		
	total sum	63673.416	102	624.249		
	benefit score sums	555.610	102	5.447		
	radical score sums	1217.167	102	11.933		
	total score sums	1761.173	102	17.266		

Table O.40 – Estimates – Idea Composite Scores

Dependent Variable	Tool	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
benefit score	A	3.179	.130	2.922	3.436
	G	3.298	.138	3.025	3.571
	V	3.179	.128	2.925	3.432
radical score	A	4.460	.192	4.080	4.841
	G	4.244	.204	3.841	4.648
	V	4.113	.189	3.737	4.488
total	A	7.639	.231	7.181	8.097
	G	7.543	.245	7.057	8.028
	V	7.292	.228	6.840	7.743
benefit sum	A	45.194	2.195	40.841	49.548
	G	44.687	2.328	40.070	49.305
	V	43.919	2.165	39.624	48.213
radical sum	A	61.917	2.953	56.060	67.774
	G	57.781	3.132	51.569	63.993
	V	57.027	2.913	51.250	62.804
total sum	A	107.111	4.164	98.852	115.371
	G	102.469	4.417	93.708	111.229
	V	100.946	4.108	92.799	109.093
benefit score sums	A	9.537	.389	8.765	10.308
	G	9.894	.413	9.076	10.713
	V	9.536	.384	8.775	10.297
radical score sums	A	13.381	.576	12.239	14.523
	G	12.733	.611	11.522	13.945
	V	12.339	.568	11.212	13.465
total score sums	A	22.918	.693	21.544	24.291
	G	22.628	.735	21.171	24.085
	V	21.875	.683	20.520	23.230

APPENDIX P

CASE STUDY – INTERVIEW GUIDES

Case Study Guides

Note: All sessions will be recorded on audiotape. This is noted on the Letter of Information and the Letter of Consent provided to the participants.

Case Study - Individual Interviewing Guide

Step 1.

Prior to the interview, provide the interviewee with the Information Letter and Letter of Consent.

Step 2.

At the start of the interview, introduce yourself to the interviewee and explain the purpose of the study and who in the organization is sponsoring this study. Ask the interviewee if they have any questions before proceeding. Also request a copy of the signed Letter of Consent. Do not proceed with the interview until the interviewee signs the Letter of Consent.

Step 3.

Proceed with the interview using the following list of questions as a guide.

Knowledge challenges

1. What knowledge challenges are you/your organization faced with? Tell me about a specific situation or challenge. And tell me about your role in the organization.

Importance of novel knowledge

2. How often does your company discover novel insights?
3. How important is it to discover novel insights, directions, ideas, and opportunities here? Why?

Discovering novel knowledge

4. How do you discover/identify novel ideas (with and/or without Athens)?
5. To what degree do tools play a role in the discovery/learning process here?
6. What other tools are currently being used to discover novel knowledge (for example, Google, Kartoo)? How are they being used?

7. Tell me about something novel that you recently discovered, specifically focusing on how you discovered it (not on the specific what – confidentiality issues).
8. Who is involved in discovering novel ideas at your organization and how do they work together? Can you provide a recent example of how this worked?
9. How often and under what circumstances do you try to find novel ideas?

Process of using an NKD tool within the organization (NKD tool from question 6)

9. How is <NKD tool> used here?
10. Do you use <NKD tool> for individual work, or group/collaborative work?
11. What motivates you to use <NKD tool>? A problem? Curiosity?
12. What types of work or problems do you specifically NOT use <NKD tool> for?
13. How do you analyze or browse through the results?
14. Do you share your <NKD tool> results with others?
15. How do you translate the <NKD tool> output into ideas?
16. How do you evaluate /decide upon which ideas to promote/act upon?

NKD tool output (NKD tool from question 6)

17. What types of knowledge do you find with <NKD tool>? Novel? Incremental?
Additional relationships? Irrelevant?
18. Does <NKD tool> provide you with knowledge that you consider beneficial?
Useful?
19. How important is this type of knowledge to the organization? To you?

Case Study - Tool Observation Sessions

Step 1.

Prior to the tool observation session, provide the participant with the Information Letter and Letter of Consent.

Step 2.

At the start of the tool observation session, introduce yourself to the participant and explain the purpose of the study and who in the organization is sponsoring this study. Ask the participant if they have any questions before proceeding. Also request a copy of the signed Letter of Consent.

Step 3.

Once the Letter of Consent has been signed, proceed with the session.

Verbal Instructions to the Participant

1. As you proceed with using the tool, tell me what you are doing and why. Also, provide me with feedback on how well (or not well) the tool is performing. In other words, tell me when you are frustrated or pleased with how the tool is performing. This will help me understand how you use the tool and why. In addition, it will help me understand where you run into challenges with either the tool's capabilities or the interpretation of the results and how the tool could be improved to alleviate these challenges.
2. I may also ask you specific questions as you proceed to better understand what you are doing and how the tool is performing for you. (Note – these cannot be pre-specified because they depend on the context of the observation session)

Feedback on Athens

1. What do you like most about Athens?
2. What challenges are you encountering when using Athens?
3. What changes do you suggest?
4. Are there specific ways you would like to use Athens, but cannot? For example, sharing results? Voting on results within your group?
5. Any other suggestions?

Case Study - Group Meeting Guide (Kick-off Meeting)

Step 1.

Prior to the group meeting, provide the participating individuals with the Information Letter and Letter of Consent.

Step 2.

At the start of the group meeting, introduce yourself to the group and explain the purpose of the study and who in the organization is sponsoring this study. Ask the group if they have any questions before proceeding. Also request copies of the signed Letter of Consent.

Step 3.

Once the Letters of Consent have all been signed, proceed with the meeting using the following agenda and list of questions as a guide.

Agenda

1. Presentation by the researcher:
 - a. Introductions around the room; discuss what they are hoping to get out of the visit.
 - b. Purpose and objectives of the visit.
 - c. Introduction to the research study.
 - d. What is novel knowledge and why is it valuable?
 - e. Introduction to Athens and training on Athens results (navigation)
 - f. Example of Athens results – how has Athens guided us towards discovering novel knowledge? (Raynaud’s syndrome query)
2. Questions to the group:
 - a. Is novel knowledge valuable to your organization? Why or why not?
 - b. Where do you look for novel knowledge? (For example, on the Web, industry publications, etc.)
 - c. What aspects of novel knowledge do you particularly value? The potential benefit or radicalness of the knowledge to the organization? To the industry?
 - d. What are the biggest challenges you face in discovering novel knowledge?

Case Study - Group Meeting Guide (Wrap-up Meeting)

Agenda

1. Summarize findings/insights thus far.
2. Include examples from an interesting search result during the visit.
3. Questions for the group
 - a. Did Athens help you discover novel knowledge?
 - b. Who will use Athens here?
 - c. How can Athens be improved to better enable you in the discovery of novel knowledge?
 - d. Are there any other insights you have gained regarding novel knowledge or Athens that you would like to share?
4. Wrap-up
 - a. Describe next steps, as outlined in the Letter of Information.
 - b. Open the discussion up for any remaining questions.
 - c. Thank everyone for his or her time.

APPENDIX Q CASE STUDY – CODES

Table Q.1 – Codes Used for Interviews (Group and Individual)

Main Code	Hierarchical Sub-Codes		
1.0 Knowledge Challenges			
2.0 Novel Knowledge	2.1 Importance		
	2.2 When and why		
	2.3 Types*		
3.0 Novel-Knowledge Discovery	3.1 Processes	3.1.1 How	
		3.1.2 Example	
		3.1.3 Who	3.1.3.1 Groups/Networks
			3.1.3.2 Requestor*
		3.1.4 Motivation	
		3.1.5 Tools	3.1.5.1 Types*
			3.1.5.2 Importance*
		3.1.6 Sharing of results	
		3.1.7 Evaluation	
	3.1.8 Translation to ideas*		
	3.2 Searching*	3.2.1 Types*	
		3.2.2 Process*	3.2.2.1 Analysis of results*
	3.3 Outcomes	3.3.1 Type of knowledge	
		3.3.2 Usefulness	
	3.4 Sources*	3.4.1 Age*	
3.4.2 Structure*			
3.4.3 Type*			
4.0 Innovation*	4.1 Approaches*		

*emergent code

Table Q.2 – Codes Used for both Interview and Tool-Use Observation

Main Code	Hierarchical Sub-Codes	
5.0 Free Nodes	5.1 Fear of new tool*	
	5.2 Iterative*	
	5.3 Role	
	5.4 Serendipity*	
	5.5 Internal knowledge*	
	5.6 Resistance to change*	
	5.7 Process - reframing*	
	5.8 Training effective*	
	5.9 Competitive intelligence application*	
	5.10 Importance of individual knowledge*	
	5.11 Ontology*	
6.0 Learning process	6.1 Context*	6.4.1 Pruning
	6.2 Goals*	
	6.3 Search-term definition	
	6.4 Information foraging	
	6.5 Intuiting	
	6.6 Interpreting	
	6.7 Group interpreting	
	6.8 Integrating	
	6.9 Institutionalizing	

*emergent code

Table Q.3 – Codes Used for Tool-Use Observation

Main Code	Hierarchical Sub-Codes	
7.0 Positive finding		
8.0 Types of problems (negative finding)	8.1 Description	8.1.1 Cluster nouns not descriptive
		8.1.2 Lack of Website description
	8.2 Display	8.2.1 Display of results more visual
		8.2.2 Focus user's attention
		8.2.3 Show connection to search term
		8.2.4 Summary
		8.2.5 Increase prominence of search terms on display
		8.2.6 Limit search results
	8.3 Errors	8.3.1 Overall
		8.3.2 Clusters empty
	8.4 Filtering Algorithm	8.4.1 Dependency on proper nouns
	8.5 Navigation	8.5.1 Inefficient organization of results
		8.5.2 Interface navigation
		8.5.3 Knowing where you are
		8.5.4 Layout confusing
		8.5.5 Understanding navigation hierarchy
		8.5.6 Refine search results to type of results
		8.5.7 Use different colours
		8.5.8 Navigation between levels of clusters
		8.5.9 Filtering and sorting
		8.5.10 Choosing level of novelty
	8.6 Quality of Web links	8.6.1 Quality of information/Lack of content
		8.6.2 Too many blogs
	8.7 Relevance Interpretation	8.7.1 Irrelevant clusters
		8.7.2 Unrelated tags
	8.8 Search	8.8.1 Choice of search terms
		8.8.2 Search iterations
		8.8.3 Search time
		8.8.4 Static search term
	8.9 Technological Frames	8.9.1 Comparing to Google
		8.9.2 Different from other tools therefore challenging to learn
		8.9.3 Expectations
8.9.4 What I was looking for		
8.9.5 Task technology fit*		
8.10 Understanding Clusters	8.10.1 Broad results	
	8.10.1 Cluster not novel*	
	8.10.2 Clusters too novel	

Main Code	Hierarchical Sub-Codes	
		8.10.3 Clusters too similar
		8.10.4 Difficult to see how results in a cluster are related
		8.10.5 Too few Web links per cluster
		8.10.6 Too many Web links per cluster
		8.10.7 Clusters too specific

*emergent code

APPENDIX R

CASE STUDY – RESULTS

Table R.1 – Case Study and Lab Experiment – Comparison of Results

Study 1 Results	Study 2 Results				
# of Sources	# of Sources	# of Comments	Problem Type	Problem Sub-type	Details
6	2	3	Description	Cluster nouns not descriptive	The nouns that describe the cluster were not helpful in understanding the contents of the cluster.
5	4	8		Lack of Website description	The Web links do not provide any description (meta-data) regarding the contents of the link. Thus, individuals don't know what the link contains unless they open it.
3	2	14	Display	Display of results more visual	A more visual display of the results is needed.
2	3	12		Focus user's attention	Including a more visual display of the results to help individuals assess which areas they want to explore in detail. In other words, help focus the individual's attention on what he or she finds most useful.
2	4	9		Show connection to search term	Desire to see how the results are either 1) connected to the search term, or, 2) how the results are themselves related or connected. Individuals want to know "why is this result here?"
1	2	9		Summary	Including a better summary of the results (and more visual) so that individuals can understand the high-level results and make a decision regarding which results to dive into more detail.
1	0			Increase prominence of search terms on display	No comments
2	0			Limit search results	No comments
1	1	1	Errors	Overall	There were some dead Web links included in the results.
2	2	2		Clusters empty	Athens did not finish processing all clusters. Thus, some clusters were empty.
4	5	26	Filtering Algorithm	Dependency on proper nouns	Problems were encountered with the clusters being

Study 1 Results	Study 2 Results				
# of Sources	# of Sources	# of Comments	Problem Type	Problem Sub-type	Details
					influenced by proper nouns (names of people and geographical places) when queries were run against the entire Web (no restriction). Pubmed queries seemed to pick up all the common Pubmed nouns (gene, genome, protein, RefSeq, Books, Genbank etc). This may be pulling the clusters in the wrong direction. This also means that the nouns beside each cluster are highly similar across clusters, and not helpful in describing the cluster.
1	2	3	Navigation	Inefficient organization of results	In order to review the results, individuals have to click on individual Web links and then go back to the novel clusters page. Individuals have to go back to the summary page if they want to move to a new set of novel clusters. This is a time consuming process.
7	4	11		Interface navigation	Overall, navigation through the tool was problematic. Individuals found it confusing and needed a lot of hand holding at first.
2	4	7		Knowing where you are	One of the interface problems was that individuals did not know where they were in the tool in terms of level (summary or novel cluster), how they got to this point (from which cluster), and what they had already looked at.
1	3	5		Layout confusing	Individuals found the overall interface, presentation of the results, and the clusters confusing.
4	4	6		Understanding navigation hierarchy	There was confusion and difficulty understanding the navigation hierarchy between the summary clusters and the novel clusters.
3	1	1		Refine search results to type of results	Allowing the user to specify the types of sources to display in the results. For example news sources, journals or scientific sources.

Study 1 Results	Study 2 Results					
	# of Sources	# of Sources	# of Comments	Problem Type	Problem Sub-type	Details
1	0				Use different colours	No comments
2	0				Navigation between levels of clusters	No comments
2	0				Filtering and sorting	No comments
1	0				Choosing level of novelty	No comments
2	4	5	Quality of Web links	Quality of information/Lack of content		Sometimes the Web link did not contain helpful information. For example, the results contained another set of search results, or just an abstract, or advertising.
1	0				Too many blogs	No comments
12	6	17	Relevance interpretation	Irrelevant clusters		Some clusters were irrelevant or unrelated to the search topic.
1	0				Unrelated tags	No comments
2	2	2	Search	Choice of search terms		Understanding how to choose an appropriate search term was challenging.
6	1	1			Search iterations	Tool doesn't easily allow for search iterations given its processing cycle.
3	2	3			Search time	Length of processing has implications for choice of search terms and search iterations. Most people are used to "instant gratification". This is related to technological frames.
2	0				Static search term	No comments
3	3	6	Technological Frames	Comparing to Google		People compared how the tool performed to Google or mentioned searching on Google.
1	4	16			Different from other tools therefore challenging to learn	Since Athens and the idea of discovering novel knowledge is so different from other tools, it makes using Athens very challenging. Individuals had difficulty "wrapping their brains around it" and changing their behaviour. Also found evidence of individuals trying to use Athens as if it were one of their search tools (using complex Boolean logic)
1	6	30			Expectations	Individuals reviewed the search results with a set of expectations regarding what they wanted to find, what they were looking for, how the results would be

Study 1 Results	Study 2 Results				
# of Sources	# of Sources	# of Comments	Problem Type	Problem Sub-type	Details
					presented and how fast the tool would search. Anything that didn't align with these expectations was viewed negatively. One participant recognized that this was the case (i.e. recognized this bias).
4	5	17		What I was looking for	Most individuals were looking for something specific in their results, rather than being open-minded and exploratory. Related to expectations.
0	6	14		Task technology fit*	The notion of task technology fit came up frequently. For example, alerts are useful for scanning, scientific content databases are suited for comprehensive searches, Google is useful for finding the latest news or a specific answer.
1	1	1	Understanding Clusters	Broad results	Cluster nouns or tags were too broad.
0	2	4		Cluster not novel*	Some cluster results were not novel – they were already known and highly related to the original search terms.
6	4	8		Clusters too novel	Some cluster results were too “out there” – too far away from the original search terms.
1	4	8		Clusters too similar	Some clusters were so similar it was difficult to detect what was different about them.
2	4	9		Difficult to see how results in a cluster are related	Difficult to see how the results of a cluster or the results between clusters were related and how they were different.
1	1	1		Too few Web links per cluster	Some clusters only contained one link.
1	2	4		Too many Web links per cluster	Some clusters contained a long list of links, which were time consuming to look through. Lack of meta-data was an issue when there were a large number of links to look through.
1	0			Clusters too specific	No comments

*emergent code

APPENDIX S

FIELD EXPERIMENT – TRAINING SCRIPT

Training Script

Slide 1 – Brief Intro

- Good morning. My name is Tracy Jenkin. I am a PhD student at Queen’s School of Business. Welcome to the training session associated with the Novel-Knowledge Discovery group experiment. I’m not sure how much background <sponsor’s name> provided you on the experiment, but it’s part of my thesis research, which focuses on novel-knowledge discovery in organizations.
- And the purpose of today’s session is to provide you with an introduction and training on Athens; and to have you provide descriptive terms for the experimental task. The reason that I need these terms today is because Athens is a batch program. So I need to pre-run your queries in Athens. I will bring the results to the experiment.
- Now before we start, I just want to make sure everyone has a copy of the information letter and consent form, which was emailed to you. People in the room have hard copies. Does anyone have any questions regarding the content?

Slide 2 – Agenda

- Here’s an overview of what we will be doing during this training session.
- We will start off with introductions around the room and on the phone. Just so I know who you all are and your role within <organization name>. And maybe why you are interested in participating in this experiment.
- Next, I will introduce you to Athens, which is the NKD tool you will be using during the experiment.
- I will introduce you to the task you will be working on for the experiment and how to choose descriptive terms to use in Athens.
- Next, I will show you how to use the results produced by Athens. I will be walking through a live example, which you can all follow along with online as well.
- Then we will all complete a survey in which you will provide me with descriptive terms for your Athens query.
- Lastly, we’ll talk about next steps.
- And of course we will have the opportunity for questions at the end, and as we go through the training session.

Slide 3 – Introductions

- Okay let’s start with introductions.

Slide 4 – Introduction to Athens

- Now, I will introduce you to Athens by telling you what Athens is and how it works.

Slide 5 – Case Example

- First, let’s discuss novel knowledge. Here’s a case example that helps illustrate novel knowledge, because it is a bit of an amorphous concept.

- There is a famous case in the knowledge discovery literature about the discovery of the connection between Raynaud's syndrome and fish oil.
- An information scientist named Don Swanson was searching for novel knowledge about Raynaud's syndrome – a condition that results in intermittent restriction of blood flow to extremities – hand and feet mostly. At the time, a cure for the condition had not been found. He did not know what to look for specifically, so he began with a review of the Raynaud's literature. During this review he was able to make a connection between Raynaud's syndrome and blood viscosity. So he reviewed the blood viscosity literature. And in this review, he found a connection to fish oil. So he reviewed the literature on dietary fish oil and found that it lowers blood viscosity. This led to the novel hypothesis that fish oil may be a useful dietary supplement to help decrease the blood viscosity in humans and therefore alleviate symptoms of Raynaud's syndrome. This hypothesis was later validated by medical researchers.
- The discovery follows a path, from the source literature, to an intermediate literature that relates both to the source and points to a target literature (that is also related to the intermediate literature). No documents contained both source and target.
- So you can see that the connection is two steps or leaps away from the original subject/topic.

Slide 6 – What is Athens

- So, what is Athens?
- Athens is a prototype tool designed to find novel knowledge. And was developed at Queen's. And I will emphasize prototype – this is not a commercial grade tool.
- By novel knowledge we mean knowledge that is not known to the organization, but that is interesting, relevant to the organization in terms of being related to what the organization does and knows, but it is related indirectly. This is what makes it novel and what makes it challenging to discover.
- So the Athens tool helps address many of the challenges associated with novel-knowledge discovery, which I will discuss next.

Slide 7 – What is Athens

- The practical problem for organizations is that it is challenging to discover novel knowledge.
- Locating novel knowledge is difficult because of information overload. Most search engines bury novel information in and amongst thousands of results. If it is on page 35 of 500, it is unlikely you will stumble upon it. It requires serendipity. Also search engines like Google rank results by global relevance – how everyone else as a whole views relevance of this page, not how it is relevant to your search and needs. It lacks understanding of your context.

Slide 8 – What is Athens

- In addition, it is difficult to recognize novel knowledge as relevant and significant because of its novelty.

Slide 9 – What is Athens

- Lastly, it is difficult to learn novel concepts that are outside of our existing mental models, and update our mental model with the novel concepts.

Slide 10 – Importance of Novel Knowledge

- So, why do we care about novel knowledge, why is it important?

- We think that it may guide organizations towards strategic opportunities that may be interesting and worth investigating.
- And via the discovery of novel knowledge and the generation of ideas, we think it may help foster innovation and accelerate and direct learning in innovative directions.
- And as a result generate competitive advantages by being a faster and better learner. First learner, learning superiority.

Slide 11 – How Athens Works

- Here's how Athens works in theory.
- First – Athens is not a traditional search tool, which brings back results that are highly related to the key terms provided – like Google.
- Using its mining algorithms, it brings back results that are two steps away, and indirectly connected to the search terms.
- If you think about the Raynaud's example, it automates the information search and foraging process that Don Swanson went through. Instead of reviewing all of that literature in the source, and intermediate literatures, Athens automates that.
- In the first step, Athens takes the key terms you provide and tries to understand what you know about those terms, using a basic search of those terms.
- Then it tries to find out what you probably already know about these terms by identifying things that are one step out from these key terms – the intermediate literature if you like. In this case, blood viscosity.
- It then tries to identify what you may not know about his subject by finding things that are related but two steps away. And this knowledge is what we refer to as novel knowledge clusters. The novel clusters discovered represent knowledge that you may not know.
- Athens helps identify interesting directions for these leaps outwards and, thus, intuition about interesting connections.

Slide 12 – How Athens Works

- Athens can also be used as a competitive intelligence tool. In this case, it tries to find out what your competitor knows about these terms, what they probably know, and what they do not know, but may be well positioned to do next.

Slide 13 – Introduction to the Task

- Now, let's talk about the task that you are going to be working on in the experiment.
- The key problem or question is: how can <organization name> generate business value from virtual worlds such as Second Life.
- What you are going to do is provide me with descriptive terms today in this meeting for me to run in Athens over the next two weeks.
- You will then use the results to generate novel ideas about how to generate business value.
- You can reuse, modify or create new ideas based on what you find in Athens.
- The Athens results are quite novel and indirectly connected, so you may see things that you don't initially understand the meaning of. But the exercise is to keep an open mind and generate novel ideas.
- Each of you will search your own results and generate up to five ideas each. I say only five because of time limits and limits on how much we can review.
- Then you will individually evaluate your own ideas. And pick one to propose forward.
- Then we will collect the ideas and you will review all of them and evaluate them all.

- In the end, we will collectively choose two ideas to propose forward.

Slide 14 – Choosing Descriptive Terms

- Now we need to discuss how to choose descriptive terms to provide to Athens. Because before we end the meeting each of you will provide me with your descriptive terms.
- The important thing to note is that the terms provided to Athens are not search terms. It's best to think of them as descriptive terms that describe an area of knowledge – the “what I already know about” area of knowledge.
- The scope of the terms is important. The terms should be of a reasonable size. Highly specific terms such as “websphere message broker” + “deployment scripting using Ant” are too specific and Athens won't find enough pages to find novel things. Whereas “internet” + “technology” are too general and will find bland novel knowledge because of limits on number of pages fetched.
- An example of ideal terms – let's say for folks working on real-time systems products, would be something like “real-time systems” + “modeling” or “model-based development”.
- Before I ask you for your terms, I will provide you with a handout that describes this in a little bit more detail

Slide 15 – Using Athens

- Now onto using Athens. First I will run through some key principles regarding the Athens output and then we will move to a live example.
- As I mentioned earlier, Athens is not a traditional search tool. It won't provide the answer to the task we are working on.
- Athens searches the Web, based on the keywords you provide, and finds knowledge that is related but indirectly connected to those search terms to help you discover novel knowledge. It will provide novel clusters that are two steps away from the original search terms. So you may see things that you don't initially understand the meaning of.
- It helps reframe your thinking about the subject so that you can come up with novel ideas. Thus keep your mind open as you review the results and avoid focusing on looking for a specific thing. This is very exploratory.
- For example, I did a search on MP3 players using “digital music” and “portable audio players”. And you would expect, in Google, for some download sites and player sites to come up. Well the results I was seeing were related to YouTube, Gaming, Wii, and DRM. These all vary in how novel they are to me (and to you). But they are a step or two away from MP3 and get you thinking about what's next for MP3.

Slide 16 – Using Athens

- The Athens output looks quite different from what you are used to seeing. I thought it would be useful to go over the organizing principles of the output.
- As we saw in the case example of novel-knowledge discovery involving Raynaud's, we start out with a subject we know about. The descriptive terms reflect this.
- Then, we find an intermediate literature – in this case clusters that are related to the key terms. And we probably know about them. Like blood viscosity.
- Athens clusters these intermediate results together.
- Next, Athens uses the intermediate clusters to find novel clusters – what we don't know about or don't know there is a connection.
- Each intermediate cluster generates one or more novel clusters.
- It's like a hierarchy – parent/child.

- Now the yellow box reflects what the Athens output contains. The intermediate clusters are listed on the Athens summary page at a very high level of detail. You can click on an intermediate cluster and be taken to the detailed novel clusters associated with it.

Slide 17 – Using Athens

- This is what the Athens results summary or home page looks like.
- Athens groups or clusters the results into two hierarchical levels. The first level is the parent level. So Athens identifies 11 major clusters of results, in this case, that relate indirectly to the search term. The terms beside each cluster are the nouns that describe that cluster.
- If you click on the cluster name or scroll down, you will see that each major cluster contains a number of additional child clusters called novel clusters.
- You can access these novel clusters by clicking on the major cluster links.
- If you scroll down you can see a summary of the novel clusters and their relationship to the main clusters.

Slide 18 – Using Athens

- Here are the novel clusters that make up major cluster 1. You can see there are nouns that describe this cluster.
- The Web links associated with each novel cluster are the Web pages Athens found and grouped together, based on similarity, to form this cluster. They relate indirectly to the search term and are described jointly by these terms.
- If you think this may be an interesting cluster to investigate to help you generate novel ideas, click on these links to see what each page is talking about.
- Note - It is helpful to click on a least a couple of the links in a cluster to get a better understanding of what the cluster is generally about.
- Feel free to jump around from novel cluster to novel cluster and between major parent clusters.
- Clicking on a link opens a new window or tab. You can always find this Athens window again. In your Web browser – Mozilla Firefox, a new tab is opened.
- If you find a broken link, just ignore and keep searching through the results.

Slide 19 – Using Athens

- You will notice descriptors at the top of the novel clusters pages. These tags are nouns that are prominent across the clusters. The larger the tag, the more prominently this tag appears across these web pages.
- If you see a tag term that interests you, you can click on it here and it will take you to the place where it shows you which novel clusters are associated with it.
- Just a different way of thinking about the novel clusters. There is no right or wrong way to explore the Athens output.
- As you can see, there are many ways of maneuvering and exploring the results.
- Some quick notes on navigating through Athens. From the summary Athens page, you can get to the novel clusters page by clicking on the novel cluster link. To get to one of the Web pages in the novel cluster, click on the Web link. Again this opens a new tab in your browser. To get back to Athens, simply click on the “clustering information” tab at the top of the browser. And to get to the Athens summary page from the novel clusters page, simply use the browser’s back button.

Show Example

- Run through the following online example. Let people play with the output and ask questions. Point out the following:
- Cluster 1
 - Novel Cluster 1 – Wii and gaming
 - Novel Cluster 3 – YouTube and copyrights
 - Novel Cluster 5 – (link 2) Pandora, webcasts and DIMA
- Cluster 2
 - Novel Cluster 2 – Photo story and DRM
 - Novel Cluster 4 – Video and movies
- Cluster 10
 - Novel Cluster 3 – Twang and Nokia (media sharing)
 - Novel Cluster 5 – Online waiter

Other notes

- There are no right or wrong answers. Some clusters won't make any sense, that's okay.
- Don't click on cluster name in novel clusters.
- If you forget where you are, go to the Athens output, and click on back button

Final Slide – Next steps

- Review instructions on choosing descriptive terms (handout). Let me know if you have questions.
- Then complete the survey.

See below for screenshots of Athens.

Clustering Information - Microsoft Internet Explorer

File Edit View Favorites Tools Help

ATHENS

Novel Information Discovery

Athens Results Summary Page

Major clusters

Root

- [Cluster 1:](#) (DiMA, YouTube, Webcasters, Media, iPhone)
- [Cluster 2:](#) (Media, Maker, Movie, Download, Sharon, Photo, Tony, Crawford, ForHome, UsersStudentsSmall)
- [Cluster 3:](#) (Media, Pocket, Smartphone, Michael, Mobile, WMP8, Internet, Hardware, Update, Allen)
- [Cluster 4:](#) (iPod, Bose, SoundDock, iPhone, Power, Dock, Multimedia, Protectors, Black, Call)
- [Cluster 5:](#) (iPod, Bose, SoundDock, Multimedia, Signup, Sound, Gift, Stores, Email, OverviewFeaturesPhotosSpecificationsAccessoriesReviews)
- [Cluster 6:](#) (iTunes, iPod, Apple, Gift, Card)
- [Cluster 7:](#) (PlaysForSure, Match, Cinema)
- [Cluster 8:](#) (iPod, Sindre, Sleek, Photo, Sony, Micro, DecportableReview, Apple, Walkman, PhotoSony)
- [Cluster 9:](#) (Neuros, CNET, Apple, Dell, iPod, Ticker, Sony, cassette, Call, iRiver)
- [Cluster 10:](#) (Mobile, Full, Global, Sony, Blu)
- [Cluster 11:](#) (Samsung, Yepp, Toddlers, Item, Beats, Zapzap, Upgrades, Firmware, Wexim, Intency)

Click on link to see novel clusters.

Summary of novel clusters

Root | [Cluster 1](#)

- Novel Cluster 1.1 Mark Hefflinger Mobile Media York
- Novel Cluster 1.2 CNET Potter Tech Digg America
- Novel Cluster 1.3 Internet Media Congress Yahoo TechCrunch
- Novel Cluster 1.4 Mark Google Tech Media Washington
- Novel Cluster 1.5 Internet Domino Lucky Yorker Vanity
- Novel Cluster 1.6 Buskirk Eliot Internet Sirius Media

Clustering Information - Microsoft Internet Explorer

File Edit View Favorites Tools Help

A T H E N S

Novel Information Discovery

generated Sun Jul 29 18:47:58 EDT 2007
"digital music" + "portable audio player"

7 clusters generated from 95 links:

*Mark Hefflinger Mobile Media York CNET Potter Tech Digg America Internet Congress Yahoo TechCrunch Google Washington
Domino Lucky Yorker Vanity Buskirk Eliot Sirius*

Cluster 1: (Mark, Hefflinger, Mobile, Media, York)

- <http://dmwmedia.com/columnists/mark-hefflinger?from=155>
- <http://www.dmwmedia.com/columnists/mark-hefflinger?from=620>
- <http://www.dmwmedia.com/columnists/mark-hefflinger?from=145>

Cluster 2: (CNET, Potter, Tech, Digg, America)

- <http://news.com.com/2100-1023-228327.html>
- http://news.com.com/Webcasters+dodge+copyright+bullet/2100-1023_3-270078.html

Cluster 3: (Internet, Media, Congress, Yahoo, TechCrunch)

- <http://www.thestandard.com/article/0,1902,28450,00.html?partner=law>
- <http://lists.essential.org/pipermail/a2k/2005-November/000768.html>
- http://news.com.com/YouTube%20erases%20clips%20per%20Japan%20media%20demand/2100-1030_3-6127857.html
- <http://www.digmedia.org/content.cfm?id=7239>
- <http://idolator.com/tunes/internet-radio/soundexchange-dima-continue-their-public-slap+fight-280161.php>
- <http://archives.cnn.com/2001/TECH/industry/08/06/radio.royalty.idg/index.html>
- <http://www.highbeam.com/doc/1G1-109959124.html>
- http://www.federatedmedia.net/tech/2007/03/tv_networks_announce_youtube_c.php
- <http://www.techlawforum.net/>
- http://www.federatedmedia.net/tech/2007/02/guardian_column_youtube_is_goo.php
- <http://www.techcrunch.com/2007/06/22/lastfm-not-joining-national-day-of-silence/>
- <http://www.techcrunch.com/2006/09/20/youtube-headed-for-good-morning-america/>
- <http://www.federatedmedia.net/tech/techcrunch/>
- <http://radiomagonline.com/news/streaming-royalty-fees-hold/>
- http://www.webcasters.org/iwa_newsletter/april.html
- http://www.webcasters.org/iwa_newsletter/may.html

Tags

Links to content in novel cluster.

Clustering Information - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Tag: Media

- [Cluster 1](#): (Mark, Hefflinger, Mobile, Media, York)
- [Cluster 3](#): (Internet, Media, Congress, Yahoo, TechCrunch)
- [Cluster 4](#): (Mark, Google, Tech, Media, Washington)
- [Cluster 6](#): (Buskirk, Eliot, Internet, Sirius, Media)

[top](#)

Tag: York

- [Cluster 1](#): (Mark, Hefflinger, Mobile, Media, York)

[top](#)

Tag: CNET

- [Cluster 2](#): (CNET, Potter, Tech, Digg, America)

[top](#)

Tag: Potter

- [Cluster 2](#): (CNET, Potter, Tech, Digg, America)

[top](#)

Tag: Tech

- [Cluster 2](#): (CNET, Potter, Tech, Digg, America)
- [Cluster 4](#): (Mark, Google, Tech, Media, Washington)

[top](#)

Tag: Digg

- [Cluster 2](#): (CNET, Potter, Tech, Digg, America)

[top](#)

Tag: America

- [Cluster 2](#): (CNET, Potter, Tech, Digg, America)

Tags

1. Athens Summary Page

ATHENS
Novel Information Discovery

Navigation

Root

- [Cluster 1: \(DiMA, YouTube, Webcasters, Media, iPhone\)](#)
- [Cluster 2: \(Media, Maker, Movie, Download, Sharon, Photo, Tony, Crawford, ForHome, UsersStudentsSmall\)](#)
- [Cluster 3: \(Media, Pocket, Smartphone, Michael, Mobile, WMPs, Internet, Hardware, Update, Allen\)](#)
- [Cluster 4: \(iPod, Bose, SoundDock, iPhone, Power, Dock, Multimedia, Protectors, Black, Call\)](#)
- [Cluster 5: \(iPod, Bose, SoundDock, Multimedia, Signup, Sound, Gift, Stores, Email, OverviewFeaturesPhotosSpecificationsAccessories\)](#)
- [Cluster 6: \(iTunes, iPod, Apple, Gift, Cards, Download, Montana, Nano, YouTube, Ratonouille\)](#)
- [Cluster 7: \(PlaysForSure, Match, CinemaNow, MusicMatch, Napster, Wal-Mart, Downloads, Pocket, Smartphones, Saltnine\)](#)
- [Cluster 8: \(iPod, Sindre, Sleek, Photo, Sony, Micro, DeportableReview, Apple, Walkman, PhotoSony\)](#)
- [Cluster 9: \(Neuros, CNET, Apple, Dell, iPod, Tricks, Sony, cassette, Cell, iRiver\)](#)
- [Cluster 10: \(Mobile, Full, Global, Sony, Blur, Aiden, Popkomm, LimeWire, Tunstall, GoTV\)](#)
- [Cluster 11: \(Samsung, Yepp, Toddlers, Item, Seats, Laptops, Upgradable, Firmware, Weight, Kitchen\)](#)

Novel Cluster 1.1 Mark Hefflinger Mobile Media York

Novel Cluster 1.2 CNET Potter Tech Digg America

Click on 'cluster link' to see novel cluster page.

2. Athens Novel Clusters Page

ATHENS
Novel Information Discovery

7 clusters generated from 95 links

[Mark Hefflinger Mobile Media York](#) [CNET Potter Tech Digg America](#) [Internet Congress Yahoo TechCrunch Google Wash](#) [Buskrk Ehot Strus](#)

Cluster 1: (Mark, Hefflinger, Mobile, Media, York)

- <http://dinnermedia.com/columnists/mark-hefflinger/?com=155>
- <http://www.dinnermedia.com/columnists/mark-hefflinger/?com=620>
- <http://www.dinnermedia.com/columnists/mark-hefflinger/?com=145>

Cluster 2: (CNET, Potter, Tech, Digg, America)

- <http://news.com/2100-1023-228327.html>
- <http://news.com/Webcasters+do+get+copyright+bullet/2100-1023-3-270078.html>

Cluster 3: (Internet, Media, Congress, Yahoo, TechCrunch)

- <http://www.thestandard.com/article/0,1902,28450,00.html?partner=law>
- <http://lists.essential.org/list/essmail/a2k2005-11/cvember000768.html>
- <http://news.com.com/YouTube%20raises%20clip%20per%20app%20demand/2100-1030-3-6127857.html>
- <http://www.dinnermedia.com/content.php?id=7239>
- <http://idolator.com/news/internet-radio/soundexchange-dma-continue-their-public-slap+fight-280161.php>
- <http://archives.cnn.com/2001/TECH/industry/08/06/radio.royalty.id/index.html>
- <http://www.hughes.com/doc/IG1-109959124.html>
- <http://www.federatedme-dia.net/tech/2007/03/iv-networks-announce-youtube-c.php>
- <http://www.techlawforum.net/>
- <http://www.federatedme-dia.net/tech/2007/02/guardian-column-youtube-is-go.php>
- <http://www.techcrunch.com/2007/06/22/lastfm-not-joining-national-day-of-silence/>
- <http://www.techcrunch.com/2006/09/20/youtube-headed-for-good-morning-america/>
- <http://www.federatedme-dia.net/tech/techcrunch/>
- <http://radioonline.com/news/streaming-royalty-fee-hold/>
- <http://www.webcasters.org/forums/letter/and.html>
- <http://www.webcasters.org/forums/letter/tray.html>

Click on 'web link' to see web page.

Or use browser BACK button to go to Athens Summary Page.

APPENDIX T
QUESTIONNAIRES USED IN THE FIELD EXPERIMENT

Novel Knowledge Experiment - Part 1

T.1 Pretest Questionnaire

Letter of Information and Informed Consent

Thank you for considering participation in my research study "Using IT to Support the Discovery of Novel Knowledge in Organizations". Before completing this questionnaire, please review the information letter provided to you during the Athens training session. This letter outlines the details of your participation. Once you have read the letter, please click on the continue button below.

Thank you for your time and interest.

Tracy Jenkin, PhD Candidate

By choosing to participate in this study, you are agreeing to the following statement: I consent to participate in the study "Using IT to Support the Discovery of Novel Knowledge in Organizations". I have read the Letter of Information, had any questions answered to my satisfaction and will keep a copy of this letter for my records. I confirm that I understand the provisions around confidentiality to protect my identity. I also understand that my participation is voluntary and I have been told who to contact if I have questions/concerns about this study.

1. Do you consent to participate in this study?

I consent to participate in this study

I do NOT consent to participate in this study

2. Please enter your first and last name:

3. Gender: Male Female

4. Current age:

5. Is English your first language? Yes No

6. Job title/description:

7. Name of your department:

Novel Knowledge Experiment - Part 1

8. Type of role in the organization:

- Senior Executive
- Director
- Functional Manager
- Project Manager
- Analyst
- Other (Please specify):

9. Number of years in current position:

10. Number of years with this organization:

11. Number of years of work experience:

12. Highest level of schooling achieved:

- High school
- College
- Bachelors
- Masters
- Doctorate
- Other (Please specify):

13. In general:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I am familiar with the topic of virtual worlds.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I am familiar with the Google search tool.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I use Google frequently.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I am familiar with the Vivisimo search tool.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I use Vivisimo frequently.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Part 1

14. In general:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I feel confident searching the Web for new knowledge.						
0	0	0	0	0	0	0
I feel confident finding information on the Web using a search engine.						
0	0	0	0	0	0	0
I feel confident reviewing the search engine results page						
0	0	0	0	0	0	0
I feel confident navigating the search engine results page by following hyperlinks.						
0	0	0	0	0	0	0

This part of the questionnaire asks you about your ability to use an unfamiliar piece of software. Often in our jobs we are told about software packages that are available to make work easier. For the following questions, imagine that you were given a new software package for some aspect of your work. It doesn't matter specifically what this software package does, only that it is intended to make your job easier and that you have never used it before.

The following questions ask you to indicate whether you could use this unfamiliar software package under a variety of conditions. For each of the conditions, please indicate your level of confidence in your ability to complete the job using the software package. Rate your confidence level by selecting a number from 1 to 10, where 1 indicates "Not at all confident", 5 indicates "Moderately confident", and 10 indicates "Totally confident".

15. Under the following conditions, I could complete the job using the software package...

Level of confidence:

(1) Not at all confident	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10) Totally confident
If there was no one around to tell me what to do as I go.									
0	0	0	0	0	0	0	0	0	0
If I had never used a package like it before.									
0	0	0	0	0	0	0	0	0	0
If I could call someone for help if I got stuck.									
0	0	0	0	0	0	0	0	0	0

Novel Knowledge Experiment - Part 1

If someone else had helped me get started.

If I had just the built-in help facility for assistance.

If I had used similar packages before this one to do the same job.

During the experiment, you will be asked to complete the search task described below. In order to complete this task you will need to choose descriptive terms, which will be provided to Athens. These terms should be generic, standard terms for the topic area you are searching. In practice, two terms is ideal. Each 'term' can be composed of one, two, or three words, creating a phrase. For example, "real-time systems" is a single term. Each term will be provided to the search tool in "quotes". For example, if real-time systems was one term and modeling the second term, the query would be "real-time systems" + "modeling". You can also use Google operators such as "+" (and) and "OR" (or) if you like. For example: "text mining OR web mining" is a single term.

Keep the scope of your terms to a reasonable size. Highly specialized terms may not provide interesting results. Terms that are too general may also be problematic and hide many interesting results. For example, terms such as "internet" + "technology" are too general; terms such as "websphere message broker" + "deployment scripting using Ant" are too specific. Terms such as "real-time systems" + "modeling" or "content management" + "text mining" work much better.

Additional instructions will be provided when the experiment starts. Please read the task and answer the questions that follow carefully.

Your Task:

The problem being addressed in this session is: How can <ORGANIZATION NAME> generate business value from virtual worlds such as Second Life? The goal of the session is to **generate novel ideas** that might lead to the creation of business value from virtual worlds for <ORGANIZATION NAME>. During the experiment, you will be asked to use the Athens search results, by browsing the clusters and web pages returned, to search for and generate these novel ideas. You will be asked to record these ideas, and then **evaluate** them on how well they address the question posed above, as well as several other measures of their potential. You will also be asked to evaluate all of the ideas generated by the group using an anonymous voting process. At the end of the session, we will **collectively choose** the top two ideas we think should be pursued by <ORGANIZATION NAME> further.

Novel Knowledge Experiment - Part 1

Below, please identify two sets of descriptive terms that you would like to use in the search tool. You will use these terms during the experimental task, so please answer thoughtfully.

As a **practical example**, a company that develops MP3 players and wants to find novel ways of applying their MP3 expertise may choose the following descriptive terms: 1) "digital music" and 2) "portable audio players".

16. Descriptive Term 1:

17. Descriptive Term 2:

Thank you for completing this questionnaire. The next step is to complete the experiment in the lab.



Ideas

T.2 Ideas Web Form

Document 4 to 5 ideas that you generate as you proceed with your search. Document the following: 1) the idea, and 2) the webpage from which you generated or found the idea.

Once you have documented an idea, click on the "next idea" button to document additional ideas.

1. Please enter your participant id:
2. At which laptop station are you sitting?

Choose One...	V
---------------	---

3. Idea 1

Description:

4. Idea 1

Web Page link:

Next Idea

5. Idea 2

Description:

6. Idea 2

Web Page link:

Next Idea

Ideas

7. Idea 3

Description:

8. Idea 3

Web Page link:

Next Idea

9. Idea 4

Description:

10. Idea 4

Web Page link:

Next Idea

11. Idea 5

Description:

12. Idea 5

Web Page link:

Done

Thank you. Please raise your hand to let me know that you have completed this task.



Novel Knowledge Experiment - Survey 2

T.3 Posttest Questionnaire – Phase 1

In the following questionnaire, you will be asked to answer a number of questions about the ideas you generated today, your experience with the Athens tool and the search task. In addition, you will be asked a few more questions about yourself.

Thank you for your time and interest.

Tracy Jenkin, PhD Candidate

Now that you have rank ordered your ideas, please take the top rated idea from your list – in other words, the best idea that addresses the search task and one that your propose <ORGANIZATION NAME> pursuer further – and answer the following questions. The confidentiality of the idea you submit will be protected by concealing your name.

1. Please briefly describe your top rated idea:
2. Please explain the reasons why the organization should pursue this idea:
3. Using the following scale, how satisfied were you with the process you used to decide which was your top rated idea?

(1) Not satisfied	(2)	(3)	(4) Moderately satisfied	(5)	(6)	(7) Highly satisfied
I was...						
0	0	0	0	0	0	0

4. Using the following scale, how confident were you with this decision?:

(1) Not confident	(2)	(3)	(4) Moderately confident	(5)	(6)	(7) Highly confident
I was...						
0	0	0	0	0	0	0

5. Please enter your participant id:

Novel Knowledge Experiment - Survey 2

Please answer the following series of questions regarding your experience with the Athens tool and the search task in today's experiment.

6. What type of knowledge did you discover with your tool?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
This tool helped me develop a broader understanding of a topic with which I am already familiar.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn incremental details about what I already know.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn concepts that are new to me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn novel associations to the topic that surprised me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn deeper knowledge about the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn relationships to other similar concepts.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me learn new concepts that I had thought were unrelated to the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me expand my knowledge of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn more specific details about what I already know.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
This tool helped me learn about new knowledge previously unknown to me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

7. How did this tool help you learn?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
The tool found the information for which I was looking.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with a summary of the results rather than all of the raw results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me to develop interesting insights from the results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with refined and meaningful results.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
With this tool, I had to refine my search terms several times in order to get the results I wanted.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Because of the interesting results it provided, the tool helped me to develop my own insights that I had not thought about previously.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool interpreted the raw results for me.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool provided me with access to the information for which I was searching.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me to develop insights about this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool required me to provide precise terminology in the search terms.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

8. How did this tool affect your understanding of this subject?

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
The tool supported my current understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool changed my perspective on this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool increased my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool made me reconsider my overall views of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me maintain my perspective regarding this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me augment my current understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me reframe my thinking regarding this concept.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me confirm my understanding of the subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool extended my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool reinforced my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool helped me refine my understanding of this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The tool re-oriented my thinking on this subject.						
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Novel Knowledge Experiment - Survey 2

9. What were you trying to accomplish in your search?

10. By completing today's task...

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I increased my familiarity with the topic of virtual worlds.						
0	0	0	0	0	0	0

11. Describe the functionality of the tool you used for today's task:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
Searches the Web using keywords entered.						
0	0	0	0	0	0	0
Groups results into clusters of similar content.						
0	0	0	0	0	0	0
Draws attention to interesting and novel results that are indirectly related to keywords entered.						
0	0	0	0	0	0	0
Allows users to drill down into clusters of similar content.						
0	0	0	0	0	0	0
Includes links to results that are most similar to keywords entered.						
0	0	0	0	0	0	0
Results are indirectly connected to keywords, rather than most similar to keywords.						
0	0	0	0	0	0	0

12. Please answer the following questions regarding the training provided:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
Sufficient training was provided at the beginning of the experiment.						
0	0	0	0	0	0	0
Additional training would have helped me complete the experimental task more effectively.						
0	0	0	0	0	0	0

Novel Knowledge Experiment - Survey 2

13. Please answer the following questions regarding the instructions provided today:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I understood what I was being asked to do in this experiment.						
0	0	0	0	0	0	0
The instructions provided by the facilitator were clear.						
0	0	0	0	0	0	0

14. Describe your satisfaction with using the Athens tool:

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
I found the tool useful for completing today's task.						
0	0	0	0	0	0	0
I found using the tool frustrating.						
0	0	0	0	0	0	0
I was better off using the tool to complete the task than without the tool.						
0	0	0	0	0	0	0

Please answer the following questions and provide examples from your search task today, where applicable:

15. What features of this tool assisted you with today's task?

16. What challenges did you encounter using this tool?

17. If this tool could be improved, what are the most important features that should be added or fixed?

Novel Knowledge Experiment - Survey 2

18. In general...

(1) Strongly Disagree	(2)	(3)	(4) Neither Agree or Disagree	(5)	(6)	(7) Strongly Agree
If I heard about a new information technology, I would look for ways to experiment with it.						
0	0	0	0	0	0	0
Among my peers, I am usually the first to try out new information technologies.						
0	0	0	0	0	0	0
In general, I am hesitant to try out new information technologies.						
0	0	0	0	0	0	0
I like to experiment with new information technologies.						
0	0	0	0	0	0	0

Thank you for completing survey 2 of this study. You have now completed phase 1 of the experiment! After a short break, we will proceed to phases 2 and 3.

Thank you!
 Tracy Jenkin, PhD Candidate



Decision Questions - Phase 2

T.4 Posttest Questionnaire – Phase 2

In the following brief questionnaire, you will be asked to answer questions about your satisfaction with the decision making process in this phase and your confidence in the decision made.

1. Using the following scale, how satisfied were you with the decision making process in this phase?

(1) Not satisfied	(2)	(3)	(4) Moderately satisfied	(5)	(6)	(7) Highly satisfied
I was...						
0	0	0	0	0	0	0

2. Using the following scale, how confident were you with this decision?:

(1) Not confident	(2)	(3)	(4) Moderately confident	(5)	(6)	(7) Highly confident
I was...						
0	0	0	0	0	0	0

Thank you for completing this questionnaire. You have now completed phase 2 of this experiment. After a short break, we will proceed to phase 3.

Thank you!
 Tracy Jenkin, PhD Candidate



Decision Questions - Phase 3

T.5 Posttest Questionnaire – Phase 3

In the following brief questionnaire, you will be asked to answer questions about your satisfaction with the decision making process in this phase and your confidence in the decision made.

1. Using the following scale, how satisfied were you with the decision making process in this phase?

(1) Not satisfied	(2)	(3)	(4) Moderately satisfied	(5)	(6)	(7) Highly satisfied
I was...						
0	0	0	0	0	0	0

2. Using the following scale, how confident were you with this decision?:

(1) Not confident	(2)	(3)	(4) Moderately confident	(5)	(6)	(7) Highly confident
I was...						
0	0	0	0	0	0	0

3. Of the three phases in this experiment, please indicate which phase you preferred most in terms of the decision making process:

- Phase 1
- Phase 2
- Phase 3

4. Please describe why you preferred this phase:

Thank you for completing this questionnaire. You have now completed phase 3 of this experiment.

Thank you!
Tracy Jenkin, PhD Candidate

Decision Questions - Phase 3



APPENDIX U

FIELD EXPERIMENT – EXPERIMENT SCRIPT

Verbal Process Guidelines:

Notes:

As people come in, introduce myself. Get them seated at computers where the lid are down and papers on top of the computer (name, participant id and list of search terms).

Slide 1 - Brief Introduction:

- Okay let's get going. Good morning.
- (Ask if anyone has not given me their consent form.)
- This study is completely voluntary so if you wish to withdraw at any time, you are free to do so. Let me know if you do not consent to video taping. In that case, I will not start the video camera.
- If you wish to participate, please sign the letter of consent and pass it to me before starting.
- As indicated, this experiment is scheduled to take 3 hours.
- Once you have all completed the consent form, we will begin.

Slide 2 - Agenda:

- Before we start the experiment, I will give you an overview of what we will be doing as well as some instructions/training review.
- First, we will take a look at the overall schedule for today's experiment.
- Next, I will give you an overview of the activities included in each phase of the experiment
- Then we will review the task you will be completing. We also discussed this in the initial training session two weeks ago.
- Next, I will provide some review training on Athens, the tool you will be using to complete your task.
- Next, you will begin your task.

Slide 3 - Schedule:

- Here's a high level overview of the schedule we will follow over the next 3 hours.
- First, phase 1 will last approximately 80 minutes.
- After completing phase 1, we will take a short 10 minute break and then proceed onto phase 2 and 3, which will last approx 45 minutes each.
- Once we are done, we will have lunch and a debriefing.
- Let's now talk about the individual phases.

Slide 4 – Phase 1:

- In phase 1 you will complete the search task previously discussed in the training session, using the Athens tool results.
- Once you have completed your task, I will give some brief training on GroupSystems.
- You will evaluate the ideas that you have generated by rating them in GroupSystems.

- Then you will answer some questions about your task, the search tool and your experience using an online survey.
- I will guide you through this process.

Slide 5 – Phase 2 and 3:

- In Phase 2, I will combine all the ideas generated in Phase 1, merging or removing duplicates.
- Then I will
 - Provide you with the list of ideas on my shared C drive (drive T)
 - You will review these ideas.
 - Then evaluate them, like you did in phase 1.
- In phase 3, we will
 - Discuss and review the results as a group.
 - Then evaluate the ideas again
 - And choose the two best ideas to propose forward.

Slide 6 - Task Intro:

- Please review the task instructions (instructions at each computer terminal in paper form). When you have finished reading raise your hand.
- Over the next 30 to 40 minutes, you will use the Athens search results to complete the task. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.
- As per the instructions, please generate up to 5 ideas.
- As you proceed in your search, document the ideas you generate from your search in the Web form “Ideas” open on your desktop. Please document the idea and the Website you were looking at which spurred you to generate the idea.
- If you are looking at the summary results page when you generate your idea, and it is the summary that prompts you to generate your idea, list that Website address in the Web page link.
- For the Web link, just copy and paste the link from the Internet Explorer address bar into this form field.

Slide 7 - Training:

- Now I will review the training on Athens

Slide 8 – Using Athens:

- Athens is a prototype tool designed to find novel knowledge. It was developed at Queen’s University. This is very different from what most search tools do. Most search tools report back to you highly related and connected terms.
- Athens searches the Web, based on the keywords you provide, and finds knowledge that is related but indirectly connected to those search terms to help you discover novel knowledge. It jumps out two leaps from those search terms to help you generate innovative, creative or novel ideas. So in many cases you may see Web links that you think “hmm I don’t see how that’s related”. What it is trying to do is pull you out there and get you thinking. View this as an exploratory process. Try not to be focused on looking for a specific thing. It is supposed to help you think in different ways and help you reframe your thinking
- As I mentioned earlier, Athens is not a traditional search tool. It won’t provide the answer to the task we are working on.

- For example, I did a search on MP3 players using “digital music” and “portable audio players”. And you would expect, in Google, for some download sites and player sites to come up. Well the results I was seeing were related to YouTube, Gaming, Wii, and DRM. Now these all vary in how novel they are to me (and to you). But they are a step or two away from MP3 and get you thinking about what’s next for MP3.

Slide 9 – Using Athens:

- In the first step, Athens takes the key terms you provide and tries to understand what you know about those terms, using a basic search of those terms.
- Then it tries to find out what you probably already know about these terms by identifying things that are one step out from these key terms – the intermediate literature if you like. For example, blood viscosity. We call these the intermediate clusters.
- It then tries to identify what you may not know about his subject by finding things that are related but two steps away. And this knowledge is what we refer to as novel knowledge clusters. The novel clusters discovered represent knowledge that you may not know.
- You can see there is a parent child hierarchy, which will be become more apparent when we look at the user interface.
- Athens helps identify interesting directions for these leaps outwards and thus intuition about interesting connections.
- Now the yellow box reflects what the Athens output contains. The intermediate clusters are listed on the Athens summary page at a very high level of detail. You can click on an intermediate cluster and be taken to the detailed novel clusters associated with it.

Slide 10 – Using Athens:

- This is what the Athens results summary or home page looks like.
- Athens groups or clusters the results into two hierarchical levels. The first level is the parent level. So Athens identifies 11 major clusters of results, in this case, that relate indirectly to the search term. The terms beside each cluster are the nouns that describe that cluster.
- If you click on the cluster name or scroll down, you will see that each major cluster contains a number of additional child clusters called novel clusters.
- You can access these novel clusters by clicking on the major cluster links.
- If you scroll down you can see a summary of the novel clusters and their relationship to the main clusters.

Slide 11 – Using Athens:

- Here are the novel clusters that make up major cluster 1. You can see there are nouns that describe this cluster.
- The Web links associated with each novel cluster are the Web pages Athens found and grouped together, based on similarity, to form this cluster. They relate indirectly to the search term and are described jointly by these terms.
- If you think this may be an interesting cluster to investigate to help you generate novel ideas, click on these links to see what each page is talking about.
- Note - It is helpful to click on a least a couple of the links in a cluster to get a better understanding of what the cluster is generally about.
- Feel free to jump around from novel cluster to novel cluster and between major parent clusters.
- Clicking on a link opens a new window or tab. You can always find this Athens window again. In your Web browser – Mozilla Firefox, a new tab is opened.

- If you find a broken link, just ignore and keep searching through the results.

Slide 12 – Using Athens:

- You will notice descriptors at the top of the novel clusters pages. These tags are nouns that are prominent across the clusters. The larger the tag, the more prominently this tag appears across these web pages.
- If you see a tag term that interests you, you can click on it here and it will take you to the place where it shows you which novel clusters are associated with it.
- Just a different way of thinking about the novel clusters. There is no right or wrong way to explore the Athens output.
- As you can see, there are many ways of maneuvering and exploring the results.
- Some quick notes on navigating through Athens. From the summary Athens page, you can get to the novel clusters page by clicking on the novel cluster link. To get to one of the Web pages in the novel cluster, click on the Web link. Again this opens a new tab in your browser. To get back to Athens, simply click on the “clustering information” tab at the top of the browser. And to get to the Athens summary page from the novel clusters page, simply use the browser’s back button.
- Are there any questions?

Slide 13 – Task Initiation

- Now please begin completing your task. I will alert you after 30 minutes to see if you are ready to move on or need more time. When you are done, please raise your hand.

Notes to Experimenter:

- At 30-minute mark, alert the group that 30 minutes has passed.
- Copy ideas into GSS from survey tool into rating. Then copy into ranking.

Slide 14 – GSS Training:

- Please close your search tool and any other open document on your desktop
- To rate the ideas you generated during your search task, we will use a software tool called GroupSystems.
- GroupSystems is primarily used for group meetings to help with brainstorming and collaboration.
- However today, we won’t be using the group capabilities of the tool, just the voting mechanisms for rating your ideas.
- To start the Group Systems software, you will double click on the colourful icon on the leftside of you desktop. Please do so now.
- You should see the following screen.

Slide 15 – GSS Training:

- When I start your session, you will be given some instructions and asked to rate your ideas according to how well the idea addresses the task. Where 1 represents “Not well at all” and 10 represents “Extremely well”.
- When you are done, click the “cast ballot button”, which is the ballot box icon up top.
- You will see the GroupSystems screen. At this point raise your hand so that I know you are ready to move to the next part.
- Next, I will transition you to the next rating session in GroupSystems.

Slide 16 – GSS Training:

- When you have finished rating, please raise your hand.
- Remember which idea you rated as your top idea. You will need this information for the survey.

Slide 17 – Survey:

- Please complete survey two on your desktop. And please make sure you enter your participant id.

Slide 18 – Phase 2 Instructions:

- Now for phase 2, I will place a handout on my shared C drive that contains a list of all of the ideas, plus the Weblink associated with that idea.
- Please review the ideas and access the Weblink if you wish.
- You will have 20 to 30 minutes to review the ideas.
- Once everyone has finished, you will then rate the ideas following the same process as in phase 1.
- Note, for the overall rating, if there is an idea that you don't understand, just skip it. I will then only transition the ideas that you were able to rate to the next rating session.
- At the end of this session, you will complete a brief questionnaire that asks you about this process.

Slide 19 – Phase 3 Instructions:

- Now in Phase 3, we will review, comment on, and discuss the results and ideas as a group for 20 to 30 minutes. The goal of the session is to pick two ideas to pursue further.
- After this discussion, you will evaluate the ideas again, just like in phase 2 – individually
- Then, as a group we will decide on which two ideas are the best after reviewing the updated rating results.
- Then we will document these two ideas in a paragraph each.

APPENDIX V

FIELD EXPERIMENT – TASK

The problem being addressed in this session is: **How can <organization name> generate business value from virtual worlds such as Second Life?** The goal of today’s session is to generate novel ideas that might lead to the creation of business value from virtual worlds for <organization name>. At the end of the session, we will collectively choose the top two ideas we think should be pursued by <organization name> further.

Please use the **Athens search results**, by browsing the clusters and web pages returned, to search for and generate novel ideas that might enable <organization name> to create business value from virtual worlds. Focus your efforts on using the Athens results to generate these ideas. Feel free to reuse, modify or create ideas based on what you find. Your objective should be to **generate up to 5 ideas** that you think are **novel** ways for <organization name> to create business value from virtual worlds. The quality of the ideas is more important than the quantity.

As you use the tool to search the Web, please **document your ideas** in the Web form (“Ideas”) open on your desktop by recording the following: 1) the **idea**, and 2) the **webpage** (webpage or Athens results page) which stimulated you to generate this idea. For example, an entry would look like the following:

At which laptop station are you sitting?

Station 1

Idea 1
Description:

Develop headphones with MP3 player embedded into them and provide a handheld remote to navigate

Idea 1
Web Page link:

www.mit.edu/hci_lab

Next Idea

After you have completed this task, you will **evaluate** each of your ideas on how well the idea addresses the question posed, as well as several other measures of its potential.

The next step will be to evaluate all of the ideas generated by the group using an anonymous voting process.