ASSESSING THE RELEVANCE OF CASES AND PRINCIPLES USING OPERATIONALIZATION TECHNIQUES

by

Bruce Martin McLaren

B.Sc. Computer Science *Cum Laude*, Millersville University, 1981

M.Sc. Computer Science, University of Pittsburgh, 1984

M.Sc. Intelligent Systems, University of Pittsburgh, 1994

Submitted to the Graduate Faculty of

the Intelligent Systems Program in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

1999

UNIVERSITY OF PITTSBURGH

FACULTY OF ARTS AND SCIENCES

This dissertation was presented by

Bruce Martin McLaren

It was defended on

December 16, 1999

and approved by

Kevin D. Ashley, Ph.D.

Committee Chairman

Martha Pollack, Ph.D.

Bruce G. Buchanan, Ph.D.

Manuela M. Veloso, Ph.D.

Copyright by Bruce Martin McLaren
1999

ASSESSING THE RELEVANCE OF CASES AND PRINCIPLES USING OPERATIONALIZATION TECHNIQUES

Bruce Martin McLaren, Ph.D.

University of Pittsburgh, 1999

Current methods in interpretive case-based reasoning (CBR) aim to enable automated reasoners to evaluate, decide, and justify arguments through reference to past cases. A limitation of these methods, however, is that they do not adequately address the gap between abstract, open-textured rules and concrete facts. This limitation is important because domains that could employ interpretive CBR, such as the law and ethics, often lack authoritative intermediate rules that can bridge this gap. Experienced human reasoners, also lacking intermediate rules, nevertheless make decisions. To the extent they are recorded, over time these decisions cumulatively bridge the abstraction gap by applying a variety of structured, yet largely implicit, heuristic *operationalization techniques*, such as instantiating principles and past cases by connecting them to critical facts and arbitrating between competing principles.

In this dissertation, I describe SIROCCO (System for Intelligent Retrieval of Operationalized Cases and COdes), an interpretive CBR system that uses operationalization techniques, discuss the language used to represent and process engineering ethics cases, and report on experiments to evaluate SIROCCO's capabilities. The experiments show that a subset of SIROCCO's operationalization techniques provide significant improvement in retrieval accuracy over several competitive methods, including a full-text retrieval system and a version of SIROCCO that does not employ operationalization techniques. SIROCCO is also capable of explaining its output, something the competitor methods are incapable of. Finally, SIROCCO is one of the first CBR systems to factor temporal considerations into similarity assessment. The experiments, however, do not demonstrate that SIROCCO's temporal knowledge improves its retrieval accuracy. Possible reasons for this preliminary result are discussed, and suggestions for improving the impact of temporal knowledge are proposed.

This work is similar to Artificial Intelligence (AI) and Law research, but it pioneers research in a domain with a less-explicit model of argumentation, and it addresses a wider range of cases than predecessor AI and Law systems. This research makes a contribution to interpretive CBR by investigating the application of abstract, open-textured rules to concrete facts, by identifying and cataloging techniques for case-based analysis and retrieval, and by testing the use of a detailed, chronological representation of the narrative description of a case.

Foreword

This has been a long and arduous journey. Along the way I've gotten married, fathered two children, become a graybeard, celebrated the conclusion of four decades of life, and almost passed into the new millennium. First, I want to thank my wife and children for their support throughout the journey. My wife, Gabriele, has been very patient, much longer than she might have been. For providing the love, support, patience, and freedom to complete this effort, I dedicate the dissertation to her. Thank you, Gabriele.

My sons, Patrick and Dominik, have provided the inspiration. My hope is that one day they will be proud of their Daddy for all of the hard work represented by this dissertation. I like to fantasize that one day my efforts will inspire the two of them in the pursuit of their own goals. Even though my work often detracted from our time spent together, they never complained and seemed to understand, in their own way, that Daddy was doing something "important" up there on the third floor.

My advisor, Kevin Ashley, has provided abundant enthusiasm, mentoring, and intellectual guidance. I have learned much from Kevin. He is a wonderful writer who continually amazes me with his ability to clearly and concisely express the importance of a concept or research idea. He is also an excellent public speaker who stresses the importance of preparation and practice in presenting one's work. Finally, I would like to thank Kevin for understanding my somewhat unusual situation as a Ph.D. student with a job and family responsibilities. He was always willing to accommodate and support the necessary juggling of time and responsibilities.

Each member of my Ph.D. committee made a significant contribution to the work. Bruce Buchanan is truly an AI visionary who managed, more than once, to get me to see the "bigger picture." Martha Pollack was a late addition to the committee but served a critical role nonetheless. Her enthusiasm was contagious, and she also helped to focus my attention on the all-important goal of actually *finishing*. Manuela Veloso, coming from a slightly different tradition of case-based reasoning, provided different and important perspectives on the work. I also want to thank Manuela for ordering some wonderful Portuguese food during ICCBR-95.

Vincent Aleven, my office-mate and *compadre* throughout the dissertation years, played a particularly important role in the development of my ideas. Vincent cheerfully (but critically) reviewed many of my conference papers, my thesis proposal, and portions of the dissertation.

Others in the Intelligent Systems Program at the University of Pittsburgh provided help along the way. Johanna Moore was very supportive and strongly encouraged me to finish the dissertation. Rich Thomason was a member of both my masters and thesis proposal committees. Steffi Brüninghaus reviewed several of my conference papers and provided valuable feedback.

I would also like to thank those who helped with knowledge engineering, the execution of the dissertation experiments, and the editing of the dissertation. Sharon Laubach, Nandita Mukhopadhyay, Bliksem Tobey, and Changwon Yoo did most of the case transcriptions for SIROCCO, my dissertation program. Elliott Evans co-designed the case-acquisition web site and contributed much of the text of Appendix B. David Allbritton provided valuable advice about the design of the experiment. Bob Jones, a statistics Ph.D. student at Carnegie Mellon University, helped design and develop the statistical model used in the experiments, and he also ran all of the nonparametric bootstrap tests. Lisa Bilecka and Jukka Keranen lent their expertise to the evaluation of SIROCCO. Roz Treger provided editorial comments for the entire dissertation.

Jeff Kaminski and Chip Bodin, managers at, respectively, the IBM Transarc Laboratory and Carnegie Group, Inc., were very flexible with my need to balance a paying job with my research. It was not easy holding down an almost full-time job throughout my dissertation years, but these two individuals helped to pave the way for me. I would also like to thank Jeff for allowing me to run the full-text retrieval methods on the Transarc hardware.

I would be remiss not to thank my parents and siblings. This is the first Ph.D. achieved by a member of my family, and it would not have been possible without their love, support, and encouragement over the years. My mother, Shirley, always encouraged my academic skills, particularly my writing. My father, Thomas, has provided a guidepost and model for my life. My sister, Beth, has always built up my confidence and challenged me to use my mind to the fullest. My other siblings, Mark and Debbie, have transformed, over the years, from blood family into personal friends. All of them provided the love.

Finally, I want to thank the Appalachian Trail. During my trek of this 2,138-mile trail from Springer Mountain, Georgia to Mount Kathadin, Maine over a six-month period in 1989, I made the fateful decision to pursue a Ph.D. In many respects, hiking the trail was similar to the pursuit of a Ph.D. One must fend off discouragement, loneliness, and fatigue to complete such a walk. So, too, with a Ph.D. One meets and befriends many fellow travelers on the trail. So, too, with a Ph.D. One feels an incredible sense of accomplishment and pride in completing such a long journey. So, too, with a Ph.D[*].

[*] The author was partially supported by a Mellon Fellowship. The research was sponsored by the National Science Foundation under grant number NSF-LIS 9720341, "Modeling Learning to Reason with Cases in Engineering Ethics: A Test Domain for Intelligent Assistance."

TABLE OF CONTENTS

1. INTRODUCTION 1

1.1. Problem Solving and Analysis in Weak Analytic Domains 5

1.2. Operationalizing Principles and Cases: An Overview of the Techniques 6

1.3. The Dissertation Project in a Nutshell 8

1.3.1. Statement of the Primary Thesis 9
1.3.2. Statement of the Secondary Thesis 10
1.3.3. An Introduction to the Computational Model: SIROCCO 10
1.3.4. A Summary of the SIROCCO Experiments 19

1.4. A Guide to the Remainder of the Dissertation 25

2. THE EXAMPLE DOMAIN: ENGINEERING ETHICS 27

2.1. What is Engineering Ethics? 27

2.1.1. Ethical Problem Solving and Decision Making in Engineering 28

2.2. The Empirical Analysis of the NSPE BER Cases 31

2.2.1. The Analysis Process and General Results 32
2.2.2. How Does the Review Board Use the Operationalization Techniques? 34
2.2.3. How Does the Review Board Use Temporal Knowledge? 45
2.2.4. Limitations of the NSPE BER Codes and Cases 50

3. SIROCCO: RETRIEVING AND ANALYZING CASES, CODES, AND OPERATIONALIZATIONS 51

3.1. SIROCCO's Architecture 51

3.1.1. An Example Target Case and Some Potentially Relevant Source Cases 53
3.1.2. Knowledge Representation in SIROCCO: An Ontology of Engineering Ethics 56
3.1.3. Stage 1: Surface Retrieval 75
3.1.4. Stage 2: Structural Mapping 85
3.1.5. The Analyzer 100

3.2. How Does SIROCCO Use the Operationalization Techniques? 108

3.3. How Does SIROCCO Use Temporal Knowledge? 111

3.4. Explanation in SIROCCO 116

4. THE SIROCCO EXPERIMENTS 118

4.1. The Experimental Design 118

4.1.1. The Top Level Experimental Approach 118

4.1.2. The Comparison Points 121
4.1.3. The Experiment Metric 121
4.1.4. Parameter Settings 126
4.1.5. The Test Cases 130
4.1.6. The Statistical Model 131

4.2. Testing the Primary Thesis 132
4.2.1. Experiment #1: Comparison of SIROCCO to Other Methods 133
4.2.2. Experiment #2: Expert Evaluation of Extra Citations 141
4.2.3. Discussion 144

4.3. Testing the Secondary Thesis 147
4.3.1. Experiment #3: SIROCCO with and without Temporal Reasoning 148
4.3.2. Discussion 149

4.4. Testing Accuracy, Efficiency, and Scalability 153
4.4.1. Experiment #4: Accuracy versus Efficiency as Case Base Size Increases 153
4.4.2. Experiment #5: Accuracy versus Efficiency as Search Increases 155
4.4.3. Experiment #6: Efficiency as Case Base Scales Up 157
4.4.4. Discussion 158

4.5. Experiment Summary 160

5. RELATED WORK 162

5.1. Related Work in Interpretive Case-Based Reasoning 162
5.1.1. TRUTH-TELLER 165
5.1.2. GREBE 168
5.1.3. BankXX 170
5.1.4. CATO 172

5.2. Related Work in Analogical Reasoning 174
5.2.1. MAC/FAC 176
5.2.2. ARCS 177

5.3. Related Work in Practical Ethics 178
5.3.1. The Ethos System 178
5.3.2. The Dax Cowart CD 179

6. CONCLUSION 181

6.1. Contributions of the Dissertation Research 182
6.1.1. The Empirical Identification of the Operationalization Techniques 182
6.1.2. The Implementation and Testing of a Computational Model That Uses Operationalization Techniques 183
6.1.3. The Development and Application of a Limited, But Expressive, Case-Representation Language 185
6.1.4. The Application of the Computational Model to a Wide Range of Cases 187
6.1.5. The Implementation and Testing of a Computational Model That Uses Temporal Knowledge to Support Retrieval and Analysis 188

6.2. Lessons Learned and Limitations of SIROCCO 189

6.3. Future Research Directions 193
6.3.1. Improving the Performance of SIROCCO 193
6.3.2. Using SIROCCO to Support Intelligent Tutoring 194

6.3.3. Developing Tools for Mapping Case Texts to the Ethics Transcription Language 195
6.3.4. Evaluating the Generality of SIROCCO and the Operationalization Techniques 196

APPENDICES 198

Appendix A: SIROCCO's Complete Knowledge Representation 199
A.1. The Ethics Transcription Language (ETL) 199
A.2. The Extended Ethics Transcription Language (EETL) 213
A.3. The Code Hierarchy 215

Appendix B: The SIROCCO Case-Acquisition Web Site 220
B.1. The Participant's Guide: An Introduction 221
B.2. Helpful Hints 222
B.3. Reading the Case 224
B.4. Creating the Transcript 226
B.5. Creating the Fact Chronology 227
B.6. Listing the Actors and objects 236
B.7. Summarizing the Board's Decision 238
B.8. Listing Cited Codes 239
B.9. Listing the Cited Cases 244
B.10. Submitting the Transcript 247
B.11. Conclusion 248

BIBLIOGRAPHY 249

References 250

INDEX 257

LIST OF FIGURES

Figure 1-1: The Operationalization Techniques employed by the NSPE BER 7
Figure 1-2: An Excerpt of SIROCCO's Output for Case 90-5-1 12
Figure 1-3: An Excerpt of SIROCCO's Explanation of the Analysis of Case 90-5-1 18
Figure 1-4: The Five Methods Compared to SIROCCO to Test the Primary Thesis 21
Figure 2-1: The Fact Situation, Question Raised, and Relevant Codes in Case 90-5-1 35
Figure 2-2: The NSPE BER Analysis of Case 90-5-1 and the Operationalization Techniques Applied 37
Figure 2-3: The Fact Situation, Question Raised, and Relevant Codes in Case 89-5-1 39
Figure 2-4: The NSPE BER Analysis of Case 89-5-1 and the Operationalization Techniques Applied 40
Figure 2-5: The Fact Situation, Questions Raised, and Relevant Codes in Case 83-1-1 42
Figure 2-6: The NSPE BER Analysis of Case 83-1-1 and the Operationalization Techniques Applied 43
Figure 2-7: The NSPE BER Analysis of Cases 83-1-2 and 83-1-3 and the Importance of Temporal Reasoning 46
Figure 2-8: The Fact Situation, Question Raised, and Relevant Code in Case 76-5-1 48
Figure 2-9: The NSPE BER Analysis of Case 76-5-1 and the Importance of Temporal Reasoning 49
Figure 3-1: SIROCCO's Architecture 52
Figure 3-2: The Example Target Case, 90-5-1 54
Figure 3-3: Three Example Source Cases, 76-4-1, 99-1-1, and 84-5-1 56
Figure 3-4: The Use of EETL and ETL and the Relationship between Fact Situations and Cases 58
Figure 3-5: The Grammar for the Ethics Transcription Language (ETL) 59
Figure 3-6: The Fact Chronology of Case 90-5-1 60
Figure 3-7: The Actors and Objects in Case 90-5-1 60
Figure 3-8: Mapping of SIROCCO's Time Qualifiers to Allen's Temporal Relations 62
Figure 3-9: Allen's Temporal Relations [Koomen, 1989, p. 3] 63
Figure 3-10: The Analysis Representation of Case 90-5-1 66
Figure 3-11: A Portion of SIROCCO's Action/Event Hierarchy 71
Figure 3-12: A Portion of SIROCCO's Code Hierarchy 72
Figure 3-13: SIROCCO's Stage 1 Algorithm 76
Figure 3-14: SIROCCO's Predefined Abstraction Levels for Matching Fact Primitives 77
Figure 3-15: The Fact Primitive and Fact Group Content Vectors for Case 90-5-1 78
Figure 3-16: Example of Target Content Vectors Retrieving Cases Using the Hash Tables 79
Figure 3-17: The Fact Primitive and Fact Group Content Vectors for Cases 76-4-1, 99-1-1, and 84-5-1 80
Figure 3-18: SIROCCO's Stage 2 Algorithm 87
Figure 3-19: The Fact Chronology of Case 76-4-1 88
Figure 3-20: The Analysis Representation of Case 76-4-1 89
Figure 3-21: The Search Tree for 76-4-1's *Code Instantiation* of III.4. and Case 90-5-1 92

Figure 3-22: Mapping of Facts at the First Ply of the Search Tree 94
Figure 3-23: SIROCCO's Analyzer Algorithm 101
Figure 3-24: An Excerpt of SIROCCO's Output for Case 90-5-1 102
Figure 3-25: The Code-Selection Heuristics 104
Figure 3-26: The Case-Selection Heuristics 106
Figure 3-27: The Additional Suggestions Heuristics 108
Figure 3-28: How the Code Operationalizations are Implemented in SIROCCO 109
Figure 3-29: How the Case Operationalizations are Implemented in SIROCCO 110
Figure 3-30: An Example Target Case, 91-6-1, and an Example Source Case, 62-7-1 112
Figure 4-1: Summary of the Experimental Design of the SIROCCO Experiments 119
Figure 4-2: Calculation of the F-Measure in the SIROCCO Experiments 122
Figure 4-3: The Key Parameters and Experiment Settings of SIROCCO 127
Figure 4-4: The Parameters and Experiment Settings of the Five Competitor Methods 129
Figure 4-5: Mean F-Measures for all Methods over all of the Trial Cases 134
Figure 4-6: Experiment #1, Paired Differences of the F-Measures for Exact Citations 137
Figure 4-7: Experiment #1, Paired Differences of the F-Measures for Inexact Citations 138
Figure 4-8: Mean F-Measures for SIROCCO Processing the Selected Topics Cases versus the Non-Selected Topics Cases 140
Figure 4-9: Experiment #2, Paired Differences of the F-Measures for Exact and Inexact Citations 143
Figure 4-10: The Fact Situation, Question Raised, and Relevant Codes (according to the NSPE BER) in Case 96-8-1 146
Figure 4-11: Experiment #3, Paired Differences of the F-Measures for Exact and Inexact Citations 149
Figure 4-12: Experiment #4, SIROCCO's Accuracy Compared to Its Efficiency as Case Base Size Increases 154
Figure 4-13: Experiment #5, SIROCCO's Accuracy Compared to its Efficiency as *N* Increases 156
Figure 4-14: Experiment #6, SIROCCO's Efficiency as Case Base Size Increases 158
Figure 5-1: An Example of TRUTH-TELLER's Case Representation [Ashley and McLaren, 1995] 167
Figure 5-2: An Example of BankXX's Case Representation [Rissland *et al.*, 1993] 171
Figure 5-3: An Example of CATO's Case Representation [Aleven, 1997] 172
Figure B-6-1: A Screen Shot of the Case-Acquisition Web Site 220

LIST OF TABLES

Table 1-1: Topics Addressed within the NSPE BER Cases and Represented in SIROCCO's Case Base 14
Table 4-1: Experiment #1, Exact Matching Comparisons Using the Nonparametric Bootstrap 134
Table 4-2: Experiment #1, Inexact Matching Comparisons Using the Nonparametric Bootstrap 135
Table 4-3: Experiment #2, Comparison of SIROCCO and EXTENDED-MG using Nonparametric Bootstrap after Expert Evaluation 143
Table 4-4: Experiment #3, Comparison of SIROCCO and NON-TEMP SIROCCO using Nonparametric Bootstrap 148