

**DO MICRO-LEVEL TUTORIAL DECISIONS
MATTER: APPLYING REINFORCEMENT
LEARNING TO INDUCE PEDAGOGICAL
TUTORIAL TACTICS**

by

Min Chi

B.S., Xi'an Jiaotong University, 1999

M.S., University of Pittsburgh, 2006

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

2009

UNIVERSITY OF PITTSBURGH
INTELLIGENT SYSTEMS PROGRAM

This dissertation was presented

by

Min Chi

It was defended on

Nov 20th 2009

and approved by

Diane Litman, Professor, Intelligent Systems Program & Department of Computer Science

Kurt VanLehn, Professor, Department of Computer Science and Engineering, ASU

Peter Brusilovsky, Associate Professor, Intelligent Systems Program & School of Library

and Information Science

Marek Druzdzal, Associate Professor, Intelligent Systems Program & School of Library and

Information Science

Jack Mostow, Professor, Robotics Institute & Department of Machine Learning

Department, CMU

Dissertation Director: Diane Litman, Professor, Intelligent Systems Program &

Department of Computer Science

DO MICRO-LEVEL TUTORIAL DECISIONS MATTER: APPLYING REINFORCEMENT LEARNING TO INDUCE PEDAGOGICAL TUTORIAL TACTICS

Min Chi, PhD

University of Pittsburgh, 2009

In this dissertation, I investigated applying a form of machine learning, reinforcement learning, to induce tutorial tactics from pre-existing data collected from real subjects. Tutorial tactics are policies as to how the tutor should select the next action when there are multiple ones available at each step. In order to investigate whether micro-level tutorial decisions would impact students' learning, we induced two sets of tutorial tactics: the "Normalized Gain" tutorial tactics were derived with the goal of enhancing the tutorial decisions that contribute to the students' learning while the "Inverse Normalized Gain" ones were derived with the goal of enhancing those decisions that contribute less or even nothing to the students' learning. The two sets of tutorial tactics were compared on real human participants. Results showed that when the contents were controlled so as to be the same, different tutorial tactics would indeed make a difference in students' learning gains. The "Normalized Gain" students out-performed their "Inverse Normalized Gain" peers. This dissertation sheds some light on how to apply reinforcement learning to induce tutorial tactics in natural language tutoring systems.

TABLE OF CONTENTS

PREFACE	xvii
1.0 INTRODUCTION	1
1.1 RESEARCH QUESTIONS	7
1.1.1 Question 1: Do Micro-level Pedagogical Tutorial Decisions Affect Students' Learning?	7
1.1.1.1 Background on Pedagogical Tutorial Tactics	7
1.1.1.2 Elicit/Tell	11
1.1.1.3 Justify/Skip-Justify	14
1.1.2 Question 2: Is Reinforcement Learning a Feasible Method to Induce Tutorial Tactics?	17
1.1.2.1 Previous research about applying RL in ITSs	17
1.1.2.2 Applying RL to Dialogue Systems vs. Natural Language Tutoring Systems	19
1.1.2.3 Whether RL Is Able To Induce Effective Tutorial Tactics Is Still An Open Question.	21
1.2 GENERAL APPROACH	24
2.0 CORDILLERA	27
2.1 STUDENT INTERFACE	28
2.2 WIZARD INTERFACE	29
2.3 AN EXAMPLE SCRIPT	30
3.0 REINFORCEMENT LEARNING PROCEDURE	34
3.1 REINFORCEMENT LEARNING FOR TUTORIAL TACTICS	35

3.2	ISSUE 1: TRAINING CORPUS	36
3.3	ISSUE 2: KNOWLEDGE COMPONENTS	37
3.3.1	Identified KCs in the Selected Domain	38
3.3.2	Tutorial Dialogue Annotation	39
3.4	KC-BASED MDPS	42
3.4.1	Issue 3: State Representation	42
3.4.1.1	Sub-Issue 1: Feature Choices	43
3.4.1.2	Sub-Issue 2: Feature Discretization	43
3.4.1.3	Sub-Issue 3: Feature Selection	43
3.4.1.4	Sub-Issue 4: Maximum Number of Features	43
3.4.2	KC-based Action	44
3.4.3	Issue 4: KC-based Reward	44
3.5	INDUCE KC-GENERAL POLICIES	44
3.6	TETREAULT AND LITMAN'S RL TOOLKIT	45
3.6.1	Expected Cumulative Reward (ECR)	46
3.6.2	Confidence Interval	47
3.6.3	An Example to Illustrate ECR and CI	48
3.7	ISSUE 5: CONFLICTING POLICIES	49
3.8	DISCUSSION	50
4.0	STUDY 1: EXPLORATORY CORPUS	52
4.1	METHODS	52
4.1.1	Participants	52
4.1.2	Materials	53
4.1.2.1	32 Knowledge Components	53
4.1.2.2	Physics Textbook	53
4.1.2.3	Pre- and Posttest	53
4.1.2.4	Domain Principles	54
4.1.2.5	Seven Training Problems	55
4.1.2.6	Training Scripts	55
4.1.2.7	Random-Cordillera	56

4.1.2.8	Human Wizards	56
4.1.2.9	Some Clarification On The Number Of KCs Appearing In This Dissertation	56
4.1.3	Procedure	56
4.1.4	Grading	57
4.1.5	Measures	58
4.2	RESULTS	60
4.2.1	Time	60
4.2.2	Learning Results	61
4.2.3	Exploratory Corpus	63
4.2.3.1	Overall Characteristics	63
4.2.3.2	KC-based Characteristics	64
4.3	DISCUSSION	66
5.0	STUDY 2: DICHOTIC GAIN (DICHGAIN) GROUP	67
5.1	APPLY RL TO INDUCE DICHGAIN POLICIES	68
5.1.1	Training Corpus	68
5.1.2	Knowledge Components	68
5.1.3	KC-based Reward	69
5.1.4	State Representation	69
5.1.4.1	Feature Choices	70
5.1.4.2	Maximum Number of Features	76
5.1.4.3	Feature Discretization	76
5.1.4.4	Feature Selection	77
5.1.5	Conflicting Policies	78
5.1.6	Summary: Procedure of Inducing Tutorial Tactics in Study 2	78
5.2	METHODS	80
5.2.1	Participants	80
5.2.2	Materials & Procedures	81
5.2.3	Grading	81
5.2.4	Measures	81

5.3	RESULTS	83
5.3.1	Compare Pre- and Post-test	83
5.3.2	Post-hoc Comparison: DichGain vs. Exploratory	84
5.3.2.1	Post-hoc Comparison: DichGain vs. Exploratory On Training Time	84
5.3.2.2	Post-hoc Comparison: DichGain vs. Exploratory On Learning Performance	85
5.3.3	Post-hoc Comparison: DichGain vs. Exploratory Tutorial Corpora	88
5.3.3.1	Post-hoc Comparison: DichGain vs. Exploratory On Overall Tutorial Decisions	88
5.3.3.2	Post-hoc Comparison: DichGain vs. Exploratory On Individual KCs	88
5.4	DISCUSSION	94
6.0	APPLYING RL TO INDUCE NORMALIZED GAIN (NORMGAIN) AND INVERSE NORMALIZED GAIN (INVNORMGAIN) TUTORING TACTICS	96
6.1	TRAINING CORPUS	97
6.2	KNOWLEDGE COMPONENTS	98
6.3	KC-BASED REWARD	99
6.4	STATE REPRESENTATION	101
6.4.1	Sub-issues 1: Feature Choices	102
6.4.1.1	Autonomy — five features	103
6.4.1.2	Temporal Situation — three features	104
6.4.1.3	Problem Solving Contextual — fifteen features	104
6.4.1.4	Performance — twelve features	107
6.4.1.5	Background — five features	109
6.4.1.6	Student Dialogue — ten features	110
6.4.1.7	Simplified Example of Deriving Fifty Features from Log Files.	111
6.4.2	Sub-issues 2: Maximum Number of Features	112

6.4.3	Sub-issues 3: Feature Discretization	112
6.4.4	Sub-issues 4: Feature Selection	114
6.4.4.1	RL-based Feature Selection	115
6.4.4.2	PCA-based Feature Selection	117
6.4.4.3	PCA and RL-based Feature Selection	120
6.4.4.4	Random Feature Selections	122
6.5	CONFLICTING POLICIES	124
6.6	SUMMARY: INDUCTION of TUTORIAL TACTICS in STUDY 3	125
6.7	INDUCED POLICIES	127
6.7.1	Source Training Corpus	128
6.7.2	Number of Features	129
6.7.3	Feature Choices	131
6.7.3.1	Autonomy Features	131
6.7.3.2	Temporal Situation Features	132
6.7.3.3	Problem Solving Contextual Features	133
6.7.3.4	Performance Features	134
6.7.3.5	Background Features	134
6.7.3.6	Student Dialogue Features	134
6.7.4	Feature Selection	136
6.8	SUMMARY: RL in STUDY 3	137
7.0	STUDY 3: NORMALIZED GAIN (NORMGAIN) VS. INVERSE NOR-	
	MALIZED GAIN (INVNORMGAIN)	139
7.1	METHODS	139
7.1.1	Participants	139
7.1.2	NormGain-Cordillera and InvNormGain-Cordillera	139
7.1.3	Materials & Procedures	140
7.1.4	Grading	140
7.1.5	Measures	140
7.2	RESULTS	141
7.2.1	Time	141

7.2.2	Learning Performance	141
7.2.2.1	Compare NormGain vs. InvNormGain conditions: Overall Learning Performance	141
7.2.2.2	Compare NormGain vs. InvNormGain Conditions: KC-based Learning Performance	143
7.2.2.3	Summary of Learning	147
7.2.3	Log Analysis	148
7.2.3.1	Overall Tutorial decision Steps	149
7.2.3.2	Comparing I-ratio Across Primary KCs	149
7.2.3.3	Comparing J-ratio Across Primary KCs	150
7.2.3.4	Summary of Log Analysis	151
7.3	DISCUSSION	151
8.0	GENERAL DISCUSSION AND CONCLUSIONS	153
8.1	POST-HOC COMPARISON	153
8.1.1	STUDY VARIATIONS	153
8.1.2	LEARNING PERFORMANCE	155
8.1.3	LEARNING PERFORMANCE ACROSS THE FOUR GROUPS	162
8.1.4	LOG ANALYSIS	163
8.1.4.1	I-Ratio	163
8.1.4.2	Justify Ratio	164
8.2	REVISITING THE TWO RESEARCH QUESTIONS	166
8.2.1	Question 1: Micro-level Pedagogical Tutorial Decisions Affect Students' Learning.	166
8.2.2	Question 2: Reinforcement Learning is a Feasible Method to Induce Tutorial Tactics.	167
9.0	CONTRIBUTIONS AND FUTURE WORK	169
9.1	CONTRIBUTION TO COGNITIVE & LEARNING SCIENCE	169
9.2	CONTRIBUTIONS TO AI&ED, ITS & EDM	172
9.3	FUTURE WORK	173
	APPENDIX A. KNOWLEDGE COMPONENTS	176

APPENDIX B. GRADING	184
B.1 GRADING PROCEDURE	184
B.2 INTER-GRADER AGREEMENT	186
APPENDIX C. BACKGROUND SURVEY	187
C.0.1 Instructions:	187
C.0.2 Questions:	187
APPENDIX D. TEXTBOOK	190
D.0.2.1 Introduction	191
D.0.2.2 Physical Quantities	192
D.0.2.3 Mass	193
D.0.2.4 Displacement	194
D.0.2.5 Velocity	195
D.0.2.6 Acceleration	196
D.0.2.7 Gravitational Acceleration	197
D.0.2.8 Force	198
D.0.2.9 Weight Force	199
D.0.2.10 Normal Force	200
D.0.2.11 Friction Force	201
D.0.2.12 Introduction to energy	202
D.0.2.13 Kinetic energy	203
D.0.2.14 Potential energy	204
D.0.2.15 Gravitational potential energy	205
D.0.2.16 Spring potential energy	206
D.0.2.17 Total mechanical energy	208
D.0.2.18 Isolated and non-isolated systems	210
D.0.2.19 Conservation of Mechanical Energy for isolated systems	211
D.0.2.20 Internal forces vs. external forces	213
D.0.2.21 Work	214
D.0.2.22 Net work	215

D.0.2.23 Conservation of Mechanical Energy for systems whose non-isolation is due to forces	216
D.0.2.24 Different choices of system	218
D.0.2.25 Conservative and non-conservative forces	219
D.0.2.26 Internal forces must be conservative	220
D.0.2.27 Summary	221
APPENDIX E. PRE- AND POSTTEST QUESTIONS.	222
APPENDIX F. PRE- AND POST-TEST KCS	236
APPENDIX G. TRAINING PROBLEMS	238
APPENDIX H. AN EXAMPLE OF STUDENT-CORDILLERA LOG FILE	241
APPENDIX I. AN EXAMPLE TUTORIAL SCRIPT	263
APPENDIX J. STUDY 2: TUTORIAL FEATURES	315
APPENDIX K. STUDY 2: DICHGAIN TUTORIAL TACTICS	317
APPENDIX L. STUDY 3: EXAMPLE LOG 50 FEATURES	329
APPENDIX M. STUDY 3: NORMGAIN AND INVNORMGAIN TUTORIAL TACTICS (FEATURES)	333
APPENDIX N. STUDY 3: NORMAGAIN TUTORIAL TACTICS (POLICIES)	336
APPENDIX O. STUDY 3: INVNORMGAIN TUTORIAL TACTICS (POLICIES)	344
BIBLIOGRAPHY	353

LIST OF TABLES

2.1	A Sample Cordillera Script	31
3.1	Example of Tutorial Dialogue with KC Labels	40
4.1	Major Principles of Work and Energy	54
4.2	Seven Training Problems	55
4.3	Exploratory Students' Learning Performance	62
4.4	Overall Characteristics On Tutorial Decisions in Exploratory Corpus	63
4.5	KC-based Exploratory Corpus	65
5.1	A Simplified Example of Part of Student Log on Training Problem P4	75
5.2	Autonomy Features Updated	76
5.3	Compare DichGain Tactics With Tutorial Tactics Under New Feature Selection Methods On Eight Primary KCs	78
5.4	DichGain Students' Pre- vs. Post-test Performance	83
5.5	DichGain vs. Exploratory Scores: Pre vs. Post-Test (No Q_{20})	86
5.6	Overall Tutorial Decision Characteristics: DichGain vs. Exploratory Corpora	89
5.7	Tutorial Decisions Per KC.	90
5.8	I-Ratio Between DichGain vs. Exploratory on a per-KC basis.	92
5.9	Justify Ratio Differences on a per-KC Basis.	93
6.1	Compare Three Corpus on Eight Primary KCs	100
6.2	An Example PCA Feature Set from the Exploratory Corpus Induced for A KC-general Tutorial Tactics	120
6.3	Issue-by-Issue Comparison of Studies 2 and 3	127
6.4	The Source Training Corpus Of the Inducing 34 Tutorial Tactics	129

6.5	The Complexity of the 34 Induced Tutorial Tactics	130
6.6	Distribution of Policy Sizes.	130
6.7	Occurrence of Autonomy Features in The Final Tutorial Tactics	131
6.8	Occurrence of Temporal Situation Features in The Final Tutorial Tactics	132
6.9	Occurrence of Problem Solving Contextual Features in The Final Tutorial Tactics	133
6.10	Occurrence of Performance Features in The Final Tutorial Tactics	135
6.11	Occurrence of Student Dialogue Features in The Final Tutorial Tactics	136
6.12	Applying 11 Feature Selection Methods to Induce 34 Tutorial Tactics	137
7.1	NormGain vs. InvNormGain on Pre- and Post-Test	144
7.2	KC-based Pre- and Post-Test Test Scores	145
7.3	Between-Group Comparison by KC-based Pre- and Post-Test Scores	146
7.3	Between-Group Comparison by KC-based Pre- and Post-Test Scores	147
7.4	Overall Characteristics of Tutorial Decisions in Exploratory Corpus	148
7.5	Compare NormGain vs. InvNormGain on I-ratio Across Eight Primary KCs	150
7.6	Compare NormGain vs. InvNormGain on J-ratio across Eight Primary KCs	151
8.1	Compare Four Groups Under the Overall Grading Criteria	157
8.1	Compare Four Groups Under the Overall Grading Criteria	158
8.1	Compare Four Groups Under the Overall Grading Criteria	159
8.2	Compare Four Groups Under the Cumulative KC-based Grading Criteria	160
8.2	Compare Four Groups Under the Cumulative KC-based Grading Criteria	161
8.2	Compare Four Groups Under the Cumulative KC-based Grading Criteria	162
8.3	Pairwise Comparison Among Four Groups On I-ratio	164
8.4	Pairwise Comparison Among Four Groups On J-ratio	165
8.5	Pairwise Comparison Among Four Groups On Number of Justification Steps	165
A1	Individual (not net) forces	177
A2	Choosing a system for COE	178
A3	Individual (not net) work	179
A4	Net work	180
A5	Individual (not net) mechanical energies	181

A6	COE, TME and isolated/non-isolated	182
A7	Kinematics	183
C1	Questions	188
C2	High School	188
C3	Advanced Placement	189
C4	College-level Math	189
C5	College-level physics	189
F1	KC Occurrences by Pre- & Post-test problem.	237
H1	An Example of Student Log on Training Problem P4	241
H1	An Example of Student Log on Training Problem P4	242
H1	An Example of Student Log on Training Problem P4	243
H1	An Example of Student Log on Training Problem P4	244
H1	An Example of Student Log on Training Problem P4	245
H1	An Example of Student Log on Training Problem P4	246
H1	An Example of Student Log on Training Problem P4	247
H1	An Example of Student Log on Training Problem P4	248
H1	An Example of Student Log on Training Problem P4	249
H1	An Example of Student Log on Training Problem P4	250
H1	An Example of Student Log on Training Problem P4	251
H1	An Example of Student Log on Training Problem P4	252
H1	An Example of Student Log on Training Problem P4	253
H1	An Example of Student Log on Training Problem P4	254
H1	An Example of Student Log on Training Problem P4	255
H1	An Example of Student Log on Training Problem P4	256
H1	An Example of Student Log on Training Problem P4	257
H1	An Example of Student Log on Training Problem P4	258
H1	An Example of Student Log on Training Problem P4	259
H1	An Example of Student Log on Training Problem P4	260
H1	An Example of Student Log on Training Problem P4	261
H1	An Example of Student Log on Training Problem P4	262

J1	Representing Sample Dialogue in Table 5.2 Using the 18 Feature Variables	
	From Study 2	316
L1	Student Autonomy Features	329
L2	Problem Solving Context	330
L3	Background Features	330
L4	Student Dialogue Features	330
L5	Temporal Features	331
L6	Performance Features	332
M1	NormGain Tutorial Tactics:	334
M2	InvNormGain Tutorial Tactics:	335

LIST OF FIGURES

1.1 A Training Problem: P4	4
1.2 An example Dialog with Cordillera	6
1.3 Elicit vs. Tell	13
1.4 Justify vs. Skip-justify	16
2.1 Student Interface	28
2.2 Students' Response Classification Window in Wizard Interface	30
3.1 ECR and CI Sample Learned policies	49
3.2 General RL Procedure For Inducing KC-based Tutorial Tactics	51
4.1 An Example of Three Grading Criteria	59
4.2 Average Time Spent Per Training Problem By the Exploratory Group	61
4.3 Learning Performance on Exploratory Group	62
5.1 The Induced Policy $\pi_{DichGain}(KC_{21}, ET)$: Gravitational Potential Energy	79
5.2 Learning Performance of Exploratory Group	84
5.3 Per Problem Time Comparison: DichGain vs. Exploratory Group	85
7.1 Compare Time Between NormGain vs InvNormGain Groups On Training Problem	142
8.1 Compare Four Groups Learning Performance under Overall Grading	155

PREFACE

First and foremost I would like to thank my advisor Kurt VanLehn. There are so many things I have learned from him. It would take a life time to thank him. A Chinese proverb says: “once a mentor, lifetime like a father,” which is what I want to express here. Kurt is not just a mentor or advisor for me. He is like a father to me especially because, like my father, every time I thought he was wrong, it turned out I was wrong. I thank him for his wisdom, sagacious guidance, unconditional support, and frequent encouragement over the years. I feel especially lucky to have worked in the strong research environment Kurt and Professor Micki Chi have built. I have met so many brilliant people over all these years and received so much help and advice from them.

I also want to give my special thanks to Professor Diane Litman. I want to thank her for agreeing to be my committee chair. Despite the disadvantage of not having Kurt around, she has taken the role of being a full advisor for me during the past year. Without her astute knowledge, guidance, great help, and wonderful information, I do not think that this year would have gone so smoothly for me. This project would not have happened without help from the whole VanLehn group and ITSPOKE group. I especially owe a lot to Dr. Pamela Jordan, who helped me to form the research questions and plans. She has been like a mentor to me for all these years. Pam and the old ITR group deserve a lot of credit for this project. I want to thank her for all the great advice and suggestions over all these years. She volunteered to help run my subjects and presented the paper for me during the summer.

I would also like to thank the rest of my committee for taking the time out of their busy lives to help me do this and for always being so willing to negotiate times, places, etc. I really appreciate their insightful comments and questions. My other committee members are Dr. Peter Brusilovsky, Dr. Marek Druzdzel, and Dr. Jack Mostow. Additionally, I want

to give my special thanks to Professor Kevin Ashely. He appears in almost every important moment in my life at University of Pittsburgh and has given me so much valuable advice and suggestions. Quite a lot of people have helped with my project in particular: Art Ward (who helped greatly when I was running subjects and writing my dissertation); Scott Silliman and Moses Hall (who helped me greatly on setting up the labs and system and more importantly, helped me to relax); Z and Linn Taylor (I feel so lucky to have them in my life; they helped me in so many aspects that it would take the whole page to list); Joel Tetreault (who helped to set up the code for this project); Bob Hausmann and Leslie Hausmann (who have given me so much help, encouragement and support); Michael Ringenberg (I enjoyed our discussion on cognitive science and learning science, and I learned greatly from him); Hua Ai (who has helped me with some of the important citations and gave me encouragements); and Xie Wei and Yuan Changhe (who have been my friends for all these years and have constantly supported my work). I also owe a lot to Kurt's other graduate students: Chad Lane, Noboru Matsuda, Chas Murray, and Stephanie Siler. I really missed our weekly graduate meetings. Even though they graduated several years ago, whenever we meet, we still feel like brothers and sisters. Thank you for all the help and support you have provided to me for all these years. There are also many people in the Intelligent System Program and Learning Research Development Center who have helped me for all these years: Prof. Janyce Wiebe, Wendy Bergstein, Patsy Guzzi, Eric Fussenegger, Jo-Anne Krevy, De Ivanhoe and many others on the second and fifth floors.

Finally, and most importantly, without my family, I most definitely would not have been able to finish this dissertation. My dearest parents and foster parents have been always so supportive while I worked toward my Ph.D. They never complained that I could not be around with them for all these years. They have always been there to support me, love me, and worry about me. I want to also say thank you to my sister, who has taken full responsibly of caring for my parents and foster parents for all these years. My parent-in-laws, Edie and Tim, who have always there for me and support me unconditionally. Kate, who has always been so much fun to talk to. She is so encouraging and sends me all kinds of funny things to make my dissertation life easier. Finally, I am dedicating this dissertation and the effort to my husband, Collin, who got no more sleep than I did for all these months and has been

my best friends for years. Meeting and marrying him is one of the most wonderful things that has happened in my life.

1.0 INTRODUCTION

Human one-on-one tutoring is one of the most effective educational interventions. Tutored students often perform significantly better than students in classroom settings [Bloom, 1984, Cohen et al., 1982]. Computer learning environments that mimic aspects of human tutors have also been highly successful. Intelligent Tutoring Systems (ITSs) have been shown to be highly effective in improving students' learning in real classrooms [Anderson et al., 1995, Koedinger et al., 1997, VanLehn, 2006].

The development of ITSs has enabled schools and universities to reach out and educate students who otherwise would be unable to take advantage of one-on-one tutoring due to cost and time constraints [Koedinger et al., 1997]. Despite the high payoffs provided by ITSs, significant barriers remain. High development costs and the challenges of knowledge engineering have prevented widespread deployment.

In order to design an effective ITS, developers must form the basic core of the system, determine what is taught, and how. Moreover, in order to increase ITSs' deployments, individual instructors should have the ability to alter the ITSs to fit their preferred teaching style and fill in with their preferred domain contents. Authoring tools [Murray et al., 2003, Alevan et al., 2006, Alevan et al., 2005, Ainsworth and Fleming, 2005] that provide support for the software-engineering aspects of development, and thus enable non-developers to implement a system, are one promising approach to this problem. These software tools allow each individual to build customized ITSs to meet his or her own needs. Generally speaking, users of authoring tools face challenges not only in developing the content to be taught, but also in determining how to interact with the students.

Most authoring tools are built with a predefined pedagogical strategy and allow domain experts to configure parameters such as the amount of help the tutor will provide. One

potential problem is that this approach assumes all students learn best using the same set of teaching strategies. However, there are no well-established domain-general pedagogical strategies in the learning and cognitive literature, and thus, the effectiveness of these predefined pedagogical strategies is often not clear. Additionally, instructors are domain experts not learning scientists. Therefore, determining how to interact with students is a challenging task for them, because they do not necessarily have a good understanding as to how these parameters will impact student performance or subsequent behaviors [Chi et al., 2004]. In order to improve their effectiveness, the authoring tools should provide more, and more effective, methods, to help instructors decide how to interact with students.

On the other hand, it is still an open question as to whether the decisions on how to interact with students would impact learning. For any form of tutoring, the tutor's behavior can be viewed as a sequential decision process wherein, at each discrete step, the tutor is responsible for selecting the next action to take. That is, the tutor's main task can be seen as deciding what action to take at each turn. Each of these tutorial decisions affects successive actions. One preferred assumption as to the effectiveness of human one-on-one tutoring has been that the human tutors are good at making such types of tutorial decisions; moreover, these decisions are responsible for students' learning gains [Chi et al., 2001, Collins and Stevens, 1982, McArthur et al., 1982, Merrill et al., 1992]. In the learning literature, the skills used to making such tutorial decisions are often referred to as pedagogical skills. More formally, Chi, Siler, and Jeong [Chi et al., 2004] define these pedagogical skills as those that "involve skillful execution of tactics, such as giving explanations and feedback, or selecting the appropriate problems or questions to ask the students."

However, little evidence has shown that either human tutors have effective pedagogical skills, or pedagogical skills were the reason the students learned. In fact, many previous studies indicated that human tutors rarely employ any pedagogical skills when tutoring [Cade et al., 2008, Chi et al., 2004, Cho et al., 2000, Core et al., 2003, Katz et al., 2007, Evens and Michael, 2006, Merrill et al., 1995, Merrill et al., 1992, VanLehn, 1999], see also [VanLehn et al., 2003]. Additionally, skillful execution of these pedagogical skills may require that tutors adapt their actions to tutorial context, which includes each student's current knowledge level and general aptitude, the subject matter under discussion, the institutional

context in which the tutoring takes place, and so on. But little evidence has been found either that human tutors are able to monitor students' understanding accurately [Chi et al., 2004], or that tutors really adapt their decisions based on the tutorial context [Putnam, 1987]. For instance, Chi, Siler, and Jeong [Chi et al., 2004] found that human tutors do not seem to process an accurate model of students' knowledge levels during the tutoring. In fact, Putnam [Putnam, 1987] found that experienced tutors did not attempt to form highly-detailed models of their students' knowledge before attempting remedial instruction; rather, each teacher appeared to move through a curriculum script to teach the individual students.

If it was not superior pedagogical skills that enabled students to learn in these previous studies [Bloom, 1984, Cohen et al., 1982], then what did cause students to learn? One indisputable explanation is instructional content and practice opportunities. For example, previous research repeatedly showed that students working with a tutor often learned significantly more than those without one [Anderson et al., 1995, Chi et al., 2008b, Koedinger et al., 1997, Lane and VanLehn, 2005, VanLehn et al., 2007a, VanLehn et al., 2005]. However, once content was controlled to be the same across all conditions, little evidence was found that there was any difference among students under different learning treatments. Several techniques have been employed to control for content. For example, in some previous studies the domain content was controlled by ensuring students worked on the same training problems with the same human tutors or on a computer tutor that was scripted by the same human tutors [Evens and Michael, 2006, VanLehn et al., 2007a, Reif and Scott, 1999]. Additionally, the content can be controlled to be equivalent by running a human tutoring condition first, videotaping the tutoring sessions, and then having another group of students watch those videotapes [Chi et al., 2008b]. Though it is often assumed that human tutors possess more effective pedagogical skills than ITSs, previous research has shown that students who were tutored under human expert tutors were no more effective than those who were tutored under ITSs [Evens and Michael, 2006, VanLehn et al., 2007a, Reif and Scott, 1999]. Therefore, it seems the large benefit of tutoring over no-tutoring found in previous studies, may be due to a difference in instructional content rather than a difference in pedagogical skills. Until recently, there have been rising doubts in cognitive science regarding the impact of pedagogical skills on students' learning [Chi et al., 2008b, Chi et al., 2004, Chi et al., 2001,

[VanLehn et al., 2007a](#)].

However, absence of evidence is not evidence of absence. The lack of evidence supporting the impact of pedagogical skills on learning does not mean these skills are irrelevant or not important. In the studies underlying this thesis, I applied and evaluated a general data-driven methodology to learn how to make these tutorial decisions from pre-existing interactivity data rather than, as is presently common, implementing a priori pedagogical theories drawn from experts.

In order to investigate the effect of pedagogical skills on learning, it was necessary to separate tutorial decisions from instructional content, strictly controlling content so that it is equivalent for all students. It is generally difficult to control tutoring content with human tutors. Computer tutors, on the other hand, permits much greater control over, and tracking of, the tutorial content than human tutors [[Evens and Michael, 2006](#), [VanLehn et al., 2007a](#), [Reif and Scott, 1999](#)]. In this thesis, Cordillera, a Natural Language (NL) tutoring system was implemented to teach college students introductory physics.

Tutoring in domains like math and science is often structured as a two-loop procedure. An outer loop selects the problem or task the student should work on next, while the inner loop governs step/level decisions during problem solving [[VanLehn, 2006](#)]. In this structure, there are two main sources of content variation: selection of different problems, and guiding students along a different solution path. In order to minimize content variation, all participants in this thesis solved the same problems and followed the same major problem-solving steps for each problem. In educational literature, the term “step” often refers to the application of a major domain principle or equation such as Newton’s Third Law of Thermodynamics during problem solving. However, in this thesis a step generally consists of multiple micro-level steps. For example, one of the training problems, P4, is defined in [Figure 1.1](#).

In order to solve P4, students need to apply several domain principles, and some principles need to be applied more than once. For instance, one of the domain principles the students need to apply to solve for P4 is the definition of Kinetic Energy ($KE = \frac{1}{2}mv^2$). More specifically, they need to apply the definition of Kinetic Energy to *the rock at T_0* . In order to do so, the tutor takes the following four **micro-level** steps: selecting the principle to apply,

A 0.6kg rock in space has a velocity of magnitude 2.0m/s at point A and kinetic energy of 7.50J at point B. What is the net work done on the rock as it moves from A to B? We define:

T_0 : the time point when the rock is at point A.

T_1 : the time point when the rock is at point B.

Figure 1.1: A Training Problem: P4

writing the corresponding equation, solving the equation, and engaging in some qualitative discussion about the principle. At each micro-level step, at least one tutorial decision needs to be made as to how to carry out the step. As a result, this study focused on pedagogical skills that govern low-level tutorial interaction decisions on a single micro-level step: for example, whether the student is told what principle to apply, or if the system elicits it from him or her with a prompt; and whether or not the student is asked to justify the answer, once he or she has completed a step.

Figure 1.2 shows a sample dialog for how student-tutor interactions would look for carrying out the four micro-steps. In Figure 1.2, each numbered line represents a dialog turn. The labels **T** and **S** designate tutor and student turns respectively. In the example shown in Figure 1.2 the tutor and student complete a single principle application, the application of the definition of Kinetic Energy to solve for KE_0 (the rock's kinetic energy at T_0) and begin to solve for v_1 (the rock's instantaneous magnitude of velocity at T_1). The tutor and student first selected a principle to apply (lines 2 and 3), wrote the corresponding equation (4), then solved it (5), and finally, discussed the principle (6-8).

Some of the tutor turns in Figure 1.2 are labeled **{ELICIT}** or **{TELL}**. This label designates a *tutorial decision step* wherein the tutor has to make a tutorial action decision deciding whether to tell the student the information or to elicit the requisite information with a question. For example, in line 2, the targeted principle is the definition of Kinetic Energy, labeled as KC_{20} , and the tutor chooses to elicit the answer from the student by asking the question, "Which principle will help you calculate the rock's kinetic energy at

1. **T:** So let's start with determining the value of KE_0 .
2. **T:** Which principle will help you calculate the rock's kinetic energy at T0? Please provide the name of the principle, not an equation. **{ELICIT}**, KC_{20}
3. **S:** Definition of kinetic energy
4. **T:** Yes, I agree. Now I will write the equation for applying the definition of kinetic energy to the rock at T0: $KE_0 = 1/2 * m * v_0^2$ **{TELL}**, KC_{20}
5. **T:** I calculate the value for $KE_0 = 0.5 * 0.6 \text{kg} * 2.0 \text{ m/s}^2 = 1.2 \text{ J}$ **{TELL}**, KC_{23}
6. **T:** We do not know the direction of the rock's velocity at T0 from the problem statement. Would the direction of v_0 affect the rock's kinetic energy at T0?
{ELICIT}, KC_{20}
7. **S:** No, it would not affect
8. **T:** Yes, I agree. This is because the kinetic energy only depends on the rock's mass and magnitude of velocity, not the direction of the rock's velocity. **{TELL}**, KC_{20}
9. **T:** We've calculated KE_0 , recall that our plan is to calculate v_1 next.

Figure 1.2: An example Dialog with Cordillera

T0? Please provide the name of the principle, not an equation.” By contrast, if the tutor elected to tell the students, then he or she would have stated, “To calculate the rock's kinetic energy at T0, let's apply the definition of Kinetic Energy.” Both steps cover the same target knowledge.

For the purposes of this thesis the term “*pedagogical tutorial tactics*” will be used to refer to the policies for selecting the tutorial action at each micro-step level when there are multiple actions available. The primary research question in this thesis is whether pedagogical tutorial tactics focused on individual interaction decisions impact students' learning.

Existing ITSs typically employ hand-coded pedagogical rules that seek to implement

existing cognitive or instructional theories. These theories may or may not have been well-evaluated. For example, in both the CTAT [Anderson et al., 1995, Koedinger et al., 1997] and Andes systems [VanLehn et al., 2005], help is provided upon request because it is assumed that students know when they need help and will only process help when they desire it. Research on gaming, however, has raised some doubts about this, by showing that students sometimes exploit these mechanisms for shallow gains thus voiding the help value [Baker et al., 2004b, Baker et al., 2004a]. It is often difficult to evaluate hand-coded rules in a tutoring system as their performance depends upon a number of factors, such as the content difficulty, the student’s incoming competence, the system’s usability, and so on. Previous researchers have largely treated the specification of tutorial tactics as a design problem: several versions of a system are created, the only difference among them being the pedagogical model employed. Data is then collected from human subjects interacting with each version of the system, and the students’ performance is then statistically compared. Due to cost limitations, typically, only a handful of alternative tutorial tactics are explored.

Recent work on ITSs has shifted focus from hand-coded tutoring designs to more data-driven methodologies. For example, ITSs researchers have used decision theory to guide the tutoring system in lieu of hand-crafted rules [Murray and VanLehn, 2006]. In this thesis, the approach adopted does not have to rely upon a priori belief about how the tutor should teach. Instead, it proposes to “learn” how to make tutorial decisions from pre-existing student-computer interactivity corpora. The machine-learning technique chosen for this task is reinforcement learning (RL). The methodology reported in this thesis is heavily motivated by previous research in non-tutoring dialog systems. In these previous studies RL has been successfully applied to improve the effectiveness of non-tutoring dialog systems [Williams et al., 2005, Walker, 2000, Singh et al., 2002]. The system employed in this thesis is a NL tutoring system named Cordillera [Jordan et al., 2007, Jordan et al., 2006]. While NL tutoring systems can be seen as complex dialogue systems, applying RL to NL tutoring systems raises certain challenges in that the research is focused on a more complex task – instruction – than most dialogue systems. Thus it is still an open question whether RL-derived policies will prove effective in an educational context. In the following paragraphs, I will describe a general methodology showing how RL was applied to derive tutorial tactics

from computer-student interactivity data. The secondary research question is: Will RL provides a feasible method to induce pedagogical tutorial tactics?

1.1 RESEARCH QUESTIONS

1.1.1 Question 1: Do Micro-level Pedagogical Tutorial Decisions Affect Students' Learning?

1.1.1.1 Background on Pedagogical Tutorial Tactics Many studies of one-on-one tutoring show that tutors tend to dominate the tutoring sessions. For instance, they take more initiative. The tutor's primary task can be seen as deciding what action to take on each turn [Chi et al., 2001, Graesser et al., 1995]. Much of this research takes an implicitly tutor-centric perspective. It assumes that the tutors' actions are primarily responsible for tutoring effectiveness based upon the way they craft and adapt their actions to the students' needs [Collins and Stevens, 1982]. Even though students can benefit from being tutored by novice tutors [Cohen et al., 1982], expert human tutors seemingly produce better learning outcomes [Lu et al., 2007, Eugenio et al., 2006]. Here both expert and novice tutors are domain experts who differ only in terms of their tutoring experience. Similar, but less significant, results were found by Chae et al. [Chae et al., 2005] and Kim, Chae and Glass [Kim et al., 2005]. In their work students' learning gains under expert tutors were larger than learning gains under novice tutors; however, their results were only marginally significant. On the other hand, it has also been shown that expert human tutors employ different tutorial tactics than novice tutors [Hume et al., 1995, Kim et al., 2005, Lu et al., 2007]. In short, these results suggest that expert tutors may be more effective than novice tutors because they make more effective tutorial decisions.

On the other hand, the majority of previous research studies have shown that human tutors may not be very effective when selecting tutorial actions and the tutors' pedagogical skills may not determine students' learning. For example, Clark, Snow, and Shavelson [Clark et al., 1976] found that human tutors' educational effectiveness was not necessar-

ily correlated with their level of training or prior experience. They conducted a comparison study between trained human tutors and rank novices in the domain of physics. Participants were tutored for five one-hour sessions that were completed in one week. Results showed that the trained tutors were no more effective than the inexperienced tutors. Chi et al. investigated three hypotheses regarding tutor effectiveness: a tutor-centered hypothesis assuming that tutoring effectiveness arises from the tutors' pedagogical skills; a student-centered hypothesis assuming it arises from the students' active generation; and an interactive hypothesis assuming that it arises from the joint effort of both the tutors and students [Chi et al., 2001, Chi et al., 2004, Chi et al., 2008b]. They found evidence supporting the latter two hypotheses, but not the tutor-centric hypothesis.

Research in computer learning environments has found a similar lack of evidence for the tutor-centric view of tutoring effectiveness. Evens and Michael conducted a series of studies comparing four learning treatments in cardiovascular physiology [Evens and Michael, 2006]. The no-tutoring condition studied a text that included examples of the correct reasoning for solving a pacemaker problem. The CIRCSIM condition solved one training problem on a tutoring system, CIRCSIM, which presented a short text passage for each incorrect step. The CIRCSIM-tutor condition solved the same training problem on a sophisticated natural language tutoring system, CIRCSIM-tutor, which replaced the text passages in CIRCSIM with typed natural language dialogue. The human tutor condition also solved the same training problem with expert human tutors. Results showed that the latter three conditions out-performed the no-tutoring condition, but the three treatments, CIRCSIM, CIRCSIM-tutor and expert human tutors, tied with each other.

While ITSs generally support students both in the selection of problems to work on and the solving of those problems, computer-aided instructional (CAI) environments generally support only the outer loop, the problem-selection loop. Previous studies have shown that when students study the same materials and solve the same problems, a CAI will be as effective as an ITS [Sleeman et al., 1989]. More recently VanLehn et al. [VanLehn et al., 2007a] compared students who studied the same material and then studied the same training problems under a variety of conditions, including expert human tutors and a variety of ITSs. All students in the study showed learning gains but no significant difference was found among

the groups. In a subsequent review of studies of human tutors, VanLehn [VanLehn, 2009] noted that human tutors were seldom more effective than moderately interactive forms of tutoring, such as step-based tutoring systems [VanLehn et al., 2007a].

In sum, previous research has suggested that tutorial content is indisputably an important source that contributes to the effectiveness of one-on-one tutoring. The effectiveness of the pedagogical tutorial tactics, however, is still an open question. In order to investigate whether pedagogical tutorial tactics *alone* will make a difference in learning, it is necessary to control such factors as the tutoring content.

In this thesis, all students studied the same subject matter, the same training problems using the same tutorial scripts, and interacted with the computer tutors using the same user interface. For each training problem, all students experienced the identical information for all of the non-tutorial decision steps, and the variance among the students was on tutorial decision steps. For any given tutorial decision step, once tutorial action was taken, the same domain content would be carried out for all students. The following example will illustrate this.

The example used here is P4 (shown in Figure 1.1), one of the seven training problems used in this dissertation. For each training problem all participants followed a two-phase strategy which consists of collaborative solution wherein the student and tutor solve the problem together (*phase 1*), followed by post-problem discussion where the student reflects upon the solution (*phase 2*).

During phase 1, the student and the tutor solve the problem together. One important characteristic of this phase is that the tutor guides the student by applying one principle at a time. For example, solving training problem P4 (shown in Figure 1.1) involves applying three major domain principles with some principles needing to be applied twice. The three domain principles are: the definition of Kinetic Energy (KE: $KE = \frac{1}{2}mv^2$), the definition of Total Mechanical Energy (TME: $TME = KE + GPE + SPE$), and the Change of Total Mechanical Energy for Non-isolated Systems ($Net_W = TME_2 - TME_1$). The solution path for P4 students followed in this dissertation was: 1) applying the definition of Kinetic Energy to solve for the rock's kinetic energy at T_0 , 2) applying the definition of Kinetic Energy to the rock's magnitude of velocity at T_1 , 3) applying the definition of Total Mechanical Energy

to solve for rock-system's Total Mechanical Energy at T_0 , 4) applying the definition of Total Mechanical Energy to solve for rock-system's Total Mechanical Energy at T_1 , and 5) applying the Change of Total Mechanical Energy for Non-isolated Systems to solve for the work done on the rock-system from T_0 to T_1 . All of the students applied one domain principle at a time and followed the same solution path in the same order of 1-5. For each domain principle application, the tutor generally would make 3-5 micro-level tutorial decisions as shown in Figure 1.2.

During phase 2 the tutor highlights the solution's main steps, reviews any confusion that students may have had during the solution, and considers how the solution varies when the problem statement is varied in certain ways. For example, in the post-problem discussion for training problem P4 (shown in Figure 1.1), the tutor would cover eight main topics, generally one topic for each domain principle. These eight topics include the discussion about whether there are any extra steps in the solution path during the problem solving, the definition of potential energy, how changing the mass of the rock would affect the final result, and so on. In this dissertation, all four groups of students went through all main topics in the post-problem discussion in the same order (shown in Appendix I). Similar to the problem solving, the difference is how these discussions were carried out. For example, in the domain of work and energy, potential energy always involves two objects, such as potential energy of block-earth pair or potential energy of block-spring pair. However, students often focus on only one object such as the block. One of the eight topics in the post-problem discussion for P4 is regarding potential energy and there were two versions of discussion: *elicit* version vs. *tell* version. An example of the *elicit* version of post-problem discussion in P4 looks like:

Tutor: In this problem, we have selected the rock as the system. Is it possible to define potential energy for the rock system?

Student: No, it is not possible.

The *tell* version of the same point is:

Tutor: In this problem, we have selected the rock as the system and we *cannot* define a potential energy for the rock system.

This project employed four types of tutorial tactics: Exploratory, Dichotic Gain (Dich-

Gain), Normalized Gain (NormGain), and Inverse Normalized Gain (InvNormGain). All four groups of participants covered the same material using the same procedure and were trained on four versions of the same tutoring system. The only difference among the different versions employed by the four groups was the policy that each tutoring system followed to make tutorial decisions.

The primary research question is: “Will pedagogical tutorial tactics focused on individual interaction decisions impact students’ learning?” For the purposes of this thesis, two types of interaction decisions, Elicit/Tell and Justify/Skip-Justify, have been made the focus of this research. For each Elicit/Tell decision step, some students would be told the information, while others would be asked to provide the information. For each Justify/Skip-Justify decision step, the tutor sometimes would execute the justification for some students and at other times the tutor would skip the justification. Each type of interaction decision is described in detail below.

1.1.1.2 Elicit/Tell During the course of one-on-one tutoring, the tutor often faces a simple decision, should he or she *elicit* the next step information from the student, or should he or she *tell* the student the next step directly? We refer to such tutorial decisions as *elicit/tell (ET) decisions*. While a lecture can be viewed as a monologue consisting of an unbroken series of tells, individual tutoring is characterized by a great deal of give and take and is thus a mixture of elicits and tells. Generally speaking, eliciting more information from the student during tutoring will result in a more interactive tutorial dialogue.

Figure 1.3 presents a pair of sample dialogues comparing elicit and tell versions of a single tutorial dialogue extracted from a log file collected in this project. Both dialogues begin and end with the same tutor turns (lines 1 and 6 in (a) and 1 and 4 in (b)). In dialogue (a) the tutor chooses to elicit twice (lines 2-3 and 4-5 respectively). In contrast, dialogue (b) covers the same domain content with two tell actions (lines 2 and 3). As a consequence, tutorial dialogue (a) is more interactive than (b).

For the purposes of this dissertation interactivity is quantified in terms of Interactivity ratio (I-ratio), which is defined as the number of elicit decisions a student received divided

(a) Elicit Version

1. **T:** So let's start with determining the value of v_1 .
2. **T:** Which principle will help you calculate the rock's instantaneous magnitude of velocity at T1? **{ELICIT}**, KC_{20}
3. **S:** definition of kinetic energy
4. **T:** Please write the equation for how the definition of kinetic energy applies to this problem at T1 **{ELICIT}**, KC_{20}
5. **S:** $ke_1 = 0.5 * m * v_1^2$
6. **T:** From $KE_1 = 0.5 * m * v_1^2$, we get $v_1^2 = KE_1 / (0.5 * m)$. We substitute KE_1 with 7.5 J and m with 0.6 kg. The equation now becomes: $v_1^2 = 7.5 \text{ J} / (0.5 * 0.6 \text{ kg})$. The magnitude of velocity is always positive, therefore, $v_1 = 5.0 \text{ m/s}$. **{TELL}**, KC_{31}

(b) Tell Version

1. **T:** So let's start with determining the value of v_1 .
2. **T:** To calculate the rock's instantaneous magnitude of velocity at T1, we will apply the definition of kinetic energy again. **{TELL}**, KC_{20}
3. **T:** Let me just write the equation for you: $KE_1 = 0.5 * m * v_1^2$. **{TELL}**, KC_{20}
4. **T:** From $KE_1 = 0.5 * m * v_1^2$, we get $v_1^2 = KE_1 / (0.5 * m)$. We substitute KE_1 with 7.5 J and m with 0.6 kg. The equation now becomes: $v_1^2 = 7.5 \text{ J} / (0.5 * 0.6 \text{ kg})$. The magnitude of velocity is always positive, therefore, $v_1 = 5.0 \text{ m/s}$. **{TELL}**, KC_{31}

Figure 1.3: Elicit vs. Tell

by the total number of ET decisions received in a given dialogue and can be expressed in:

$$\mathbf{I - ratio} = \frac{N_{Elicit}}{N_{Elicit} + N_{Tell}} \quad (1.1)$$

The higher this value, the more interactive the dialogue. If $I - ratio = 0.5$ means that students were given elicitation prompts as often as they were simply told the information while $I - ratio > 0.5$ means that they were more likely to be prompted for information, and $I - ratio < 0.5$ means that the conversation was more didactic.

A key characteristic of one-on-one tutoring, whether from human tutors or computer support, is high interactivity. A common assumption, often referred as the *monotonic interaction hypothesis* [VanLehn et al., 2007a] is that greater interactivity leads to greater learning. But Chi et al. [Chi et al., 2001, Chi et al., 2008b] and Rose [Rose et al., 2001] found no difference in learning between students tutored on an interactive tutor and those tutored on a more didactic one. A detailed review of the literature [VanLehn et al., 2005, VanLehn, 2009] (submitted) distinguished between the widely-accepted, monotonic interactivity hypotheses and the better supported *interaction plateau hypothesis*. The former states that an increase in interactivity causes consistent increases in learning gains, while the latter states that beyond a given threshold point, increasing interactivity will yield diminishing educational returns. In this dissertation, it will be argued that it may not be the absolute volume of interactivity that is at issue, but rather how the interactivity is guided.

Some existing theories of learning suggest that when deciding whether to elicit or tell, a tutor should take into account several factors including the students' current knowledge model. Vygotsky [Vygotsky, 1971] coined the term "zone of proximal development" (ZPD) to describe the space between abilities that a student may display independently and those that they may display with support. He hypothesized that the most learning occurs when students are assigned tasks within their ZPD. In other words, the task should neither be so simple that students can achieve it independently or trivially, nor so difficult that they simply cannot make progress even with assistance. We expect, based upon this theory, that if students are somewhat competent at a given step, the tutor should elicit, and provide help only if the students fail, so that they can practice their knowledge. If students are completely unfamiliar with the step, however, then the tutor should tell them directly. Collins, Brown and Newman [Collins et al., 1989] describe a progression from tells to elicits following their "model, scaffold & fade" rubric. Koedinger and Alevan [Koedinger and Alevan, 2007] by contrast defined an "assistance dimension", which includes elicits and tells. The level of

assistance a tutor should provide may be resolved differently for different students and should be adapted to: the learning environment, the domain materials used, the students' knowledge level, their affect state, and so on.

1.1.1.3 Justify/Skip-Justify The second tutorial decision investigated was to execute or to skip a justification step. During the tutoring process, human tutors sometimes ask students to *justify* a step they have taken or an entry they have made. We refer to such tutorial decisions as *justify/skip-justify (JS) decisions*. Their apparent goal appears to be to help students understand domain knowledge in a deeper way. The open question is when should tutors conduct an elaborate discussion of a problem solving step when this discussion is not necessary for the solution? Some authors including [Chi et al., 1994], [Conati and VanLehn, 2000], [Aleven et al., 2004] and others have found that asking students to justify their solution steps improves learning. However, eliciting such a discussion may not always be desirable if, for example, the student is well aware of the rationale. If so, typing in a justification can be slow, frustrating, and distracting. Katz, O'Donnell, and Kay [Katz et al., 2000] found that in some cases it may be better to delay discussion of the justifications until after the problem has been solved, especially if the justification is abstract, plan-based, or lengthy.

After a JS decision is made and the tutor has decided to execute a justification step, the tutor sometimes needs to make an ET decision immediately. Thus, there are three possible decisions for these decision steps. Figure 1.4 presents three dialogue examples. Among them, parts in (a) and (b) justification is employed to guide the student. More specifically, in parts (a) and (b), the tutor first made a JS decision and decided to execute the justification step. Then the tutor needed to make another decision about whether to elicit the information from the students or to tell the students the justification directly. Examples (a) and (b) shows what a tutorial dialogue may look like if the tutor decides to execute an elicit justification and a tell justification respectively. In part (c), the justification is skipped. As before, all three examples start and end with the same tutor turns (lines 1 and 5 in part (a) and lines 1 and 4 in part (b) and (c)). In part (a) the tutor asked the student to justify his or her answer (lines 3-4); in part (b) the tutor told the student the justification directly (line 3),

while in part (c) the tutor skipped the justification (line 3).

Similar to I-ratio, justification can be quantified in terms of Justification ratio (J-ratio), which is defined as the number of times the tutor does not skip a justification step divided by the total number of JS decisions the tutor made and can be expressed in:

$$\mathbf{J - ratio} = \frac{N_{Justify}}{N_{Justify} + N_{SkipJustify}} \quad (1.2)$$

The higher this value, the deeper and more elaborate the dialogue might be. In addition to J-ratio, justification can also be quantified by the number of justification steps students receive. The more justification steps in a student's tutorial dialogue, the more deep learning the student may gain.

Neither the ET (elicit/tell) nor JS (justify/skip-justify) steps are well-understood, and there is no widespread consensus on how or when they should be used. This is why one research objective is to derive tutorial tactics for them from empirical data. Given that expert human tutors do not always make optimal, or effective, tutorial decisions, hand-crafting rules based upon human data may be a suboptimal strategy. Inducing data directly from interaction data may be the most suitable method.

(a) Justify Version & Elicit

1. **T:**Can we infer the direction of the velocity of the rock at T1 from the rock's kinetic energy at T1? **{ELICIT}**, KC_{20}
2. **S:**nope.
3. **T:**Excellent! Please explain why. **{JUSTIFY, ELICIT}**, KC_{20}
4. **S:**Only the magnitude of the velocity and not the direction of it is part of the definition of kinetic energy.
5. **T:**Excellent! Now that we know v_1 , we will concentrate on the sought quantity for this problem, that is, the net work done on the rock system from A to B.

(b) Justify Version & TELL

1. **T:**Can we infer the direction of the velocity of the rock at T1 from the rock's kinetic energy at T1? **{ELICIT}**, KC_{20}
2. **S:**nope.
3. **T:**Excellent! This is because the kinetic energy only depends on mass and the magnitude of velocity, not the direction of velocity.**{JUSTIFY, TELL}**, KC_{20}
4. **T:**Now that we know v_1 , we will concentrate on the sought quantity for this problem, that is, the net work done on the rock system from A to B.

(c) Skip-justify Version

1. **T:**Can we infer the direction of the velocity of the rock at T1 from the rock's kinetic energy at T1? **{ELICIT}**, KC_{20}
2. **S:**no, we cannot
3. **T:**Excellent! **{Skip-JUSTIFY}**, KC_{20}
4. **T:**Now that we know v_1 , we will concentrate on the sought quantity for this problem, that is, the net work done on the rock system from A to B.

Figure 1.4: Justify vs. Skip-justify

1.1.2 Question 2: Is Reinforcement Learning a Feasible Method to Induce Tutorial Tactics?

1.1.2.1 Previous research about applying RL in ITSs RL has been applied to conventional ITSs and used successfully to improve system behavior. In [Beck et al., 2000] the authors applied RL to develop a tutorial policy that would minimize the time students take to complete a problem. In Beck et al.’s study, the resulting policy caused the students to spend significantly less time per problem than their peers who did not follow the policy. However, the authors used simulated data for the training datasets, making it possible to accurately model time on task. As a consequence, they faced no problems of data sparsity. In this dissertation, given that the cause of student learning is still an open question, it would be difficult to accurately simulate students’ responses to the tutor and simulate how students would learn. Therefore, we used a training corpus collected from real human subjects and, due to the high cost of collecting educational data, the training corpus is comparatively small.

Barnes and Stamper [Barnes and Stamper, 2008, Stamper et al., 2007] have applied RL to automatically construct problem solutions for an ITS called Proofs Tutorial, which teaches college-level discrete mathematics. In their work, each student’s solution is defined as a diagraph with a series of states connected by actions. A state is represented by the list of premises generated in the solution so far and the actions are axiom, principle, or rules applications taken at each step. The authors collected and merged all of the student solutions into a single super-graph by taking the union of all possible states and transitions. Once the super-graph was constructed, it included all previously examined paths taken by students in solving the problem. The authors then applied MDP to induce an optimal solution using the super-graph as the search space. More specifically, they assigned scores of 100 to the goal state and -10 to each incorrect state, and a cost of 1 for each action taken. They then used value iteration to calculate the value for each state in this single graph, which was then used to generate hints for new students. They found that the extracted MDPs and the proposed hint-generating functions were able to provide hints over 80% of the time.

The work described in [Tetreault et al., 2007, Tetreault and Litman, 2006a] used RL to

learn tutorial tactics governing whether or not the system should provide feedback and what type of questions it should ask. They used a previously collected corpus of physics tutorial dialogues. Their state representation consisted of five feature vectors: *Certainty*, *Correctness*, *Percent Correct*, *Concept Repetition*, and *Frustration*; and they defined four possible tutor actions: ask a simple answer question, a complex answer question, a combination of the two, or none at all. As with the present work, they used students' final normalized learning gains as reward. Their work is close to that of this thesis because both seek to apply RL to derive effective pedagogical tutorial tactics. However, their dataset was not collected with the goal of exploring the full range of tutorial tactics in that the tutor often executed only one type of action in many dialogue states. Additionally, manually annotated features such as *Certainty* and *Frustration* were used in their work while only features that could be computed automatically or evaluated objectively, such as gender, were included in this thesis and their feature space is also substantially smaller than the one explored in this thesis. Moreover, the learned policies in their work were not tested on real students and thus their predicted success was not verified empirically. In this dissertation, the research is based on both an exploratory corpus designed to test the range of tutor actions and conducting empirical tests of the resulting KC-based strategies using real human subjects.

In this dissertation, RL was applied to induce pedagogical tutorial tactics from student-computer interactivity data. The studies tackled two challenges: the high cost of collecting a training corpus and the lack of prior knowledge as to what information to include in the state representations. Previous research on applying RL in ITSs focused on some, but not both, issues addressed here. While there have been other methods for deriving effective pedagogical tutorial tactics, they have typically involved other machine-learning methods.

Murray and VanLehn [Murray and VanLehn, 2006], for instance, applied decision theory to determine the type of hints and feedbacks the tutor should give. In their work, a Dynamic Decision Network was applied in an ITS that would decide the best actions. Their approach showed that human tutors would agree with the ITS's actions more frequently than a random tutor. However, a Dynamic Decision Network requires knowing the utility function for each state and deriving it is not trivial. The goal of this thesis research is to improve students' learning gains and thus the utility function is only available for the last state. A Dynamic

Decision Network would be required to perform look-ahead search, like a chess program, all the way to the end before it could select a single move. Consequently, it would not be straightforward to apply Dynamic Decision Networks to the current research.

As mentioned before, this research is chiefly motivated by the previous work on applying RL to non-tutoring dialogue systems. However, Natural Language (NL) tutoring systems differ from the non-tutoring dialogue system and thus it is not clear that RL is a feasible approach here.

1.1.2.2 Applying RL to Dialogue Systems vs. Natural Language Tutoring Systems

Dialogue Systems is a field of Computer Science that focuses on the construction of computer systems that interact with human users via natural-language dialogues. Much of the work in this area is focused on systems that obtain information or search databases such as querying bus schedules [Raux et al., 2005], booking airline tickets [Rudnicky et al., 1999], and accessing train schedules [Swerts et al., 2000]. NL tutoring systems can be seen as systems that use natural dialogue for instructional purposes such as helping students to learn a subject by engaging in a natural language conversation. Auto-tutor [Graesser et al., 2001], WHY-Atlas [VanLehn et al., 2002], and ITSPOKE [Litman and Silliman, 2004] for example, are all NL tutoring systems that teach students conceptual physics. For both general dialogue systems and NL tutoring systems the central component is the dialogue manager. At each point in the dialogue, it decides which action to take. In recent years, work on the design of dialogue systems has involved an increasing number of data-driven methodologies. Among these, Reinforcement Learning (RL) has been widely applied [Williams et al., 2005, Walker, 2000, Singh et al., 2002].

RL is a machine learning method that centers on the maximization of expected rewards. RL has many features well-suited to the problem of designing the dialogue manager such as unobservable states, delayed rewards, and so on. Its primary advantage is its ability to compute an optimal policy within a much larger search space, using a relatively small training corpus. It is data-efficient because it evaluates actions as a function of states.

Much of the previous research on the use of RL to improve dialogue systems such as [Levin and Pieraccini, 1997, Singh et al., 1999] has used Markov Decision Processes (MDPs)

[Sutton and Barto, 1998] to model the dialogue data and then optimize the policies from the training corpus. An MDP formally corresponds to a 4-tuple (S, A, T, R) , in which: $S = \{S_1, \dots, S_n\}$ is a state space; $A = \{A_1, \dots, A_m\}$ is an action space represented by a set of action variables; T is a set of transition probabilities between states that describe the dynamics of the modeled system; and $R = r(s_i, s_j, a_k)$ denotes a reward model that assigns rewards to state transitions and models payoffs associated with such transitions. The goal of RL is to find an optimal policy π^* that maps each state to the proper actions that would generate the maximum rewards. The dialogue management problem can be naturally cast into the MDP formalism: the states $\{S_1, \dots, S_n\}$ in the MDP correspond to the dialogue states (or an abstraction thereof); the actions $\{A_1, \dots, A_m\}$ correspond to the particular actions the dialogue manager might take; and the rewards $r(s_i, s_j, a_k)$ are defined to reflect a dialog performance metric, such as learning gains. Once the MDP structure has been defined, the transition probabilities between states T are estimated from a training corpus of dialogues, and, based on them, the policy which maximizes the expected cumulative reward is computed.

An MDP describes a stochastic control process and the state transitions possess the Markov property, which assumes that only the present state gives any information about the future behavior of a process and knowledge of the history of a process does not add any new information. However, in many real-world applications, including this thesis, the Markov property does not always hold. For example, in order to construct the MDP model, one needs to simplify the sample dialogue data, reducing it to a set of computable features. In doing so, we cannot avoid losing information some of which may be relevant. Given that these defined features do not represent the whole state, the state representations in this research do not have the Markov property. However, previous studies have shown some theories which apply to cases exhibiting the Markov property, can also be applied to many problems that do not, such as the task domain of this thesis [Williams et al., 2005, Walker, 2000, Singh et al., 1999, Singh et al., 2002]. This is because the induced policies may not be optimal, but they can still improve the effectiveness of the system and in most of these previous studies, the baseline policy is making random decisions.

While most previous work on using MDPs to train dialogue systems has been successful

[Walker, 2000, Henderson et al., 2005], whether it can be used to improve the effectiveness of NL tutoring systems is still an open question. One major source of uncertainty comes from the fact that the rewards used in RL are much more delayed in NL tutoring systems than those in non-tutoring dialogue systems. Non-tutoring dialogue systems often use user satisfaction or task completion as the rewards, while the most preferable rewards for NL tutoring systems are students' learning gains. More immediate rewards are more effective than more delayed rewards for RL induction. This is because the issue of assigning credit for a decision, attributing responsibility to the relevant decision is substantially easier in the former case. The more we delay success measures from individual decisions, the more difficult it becomes to identify the decision(s) responsible for our success or failure. Even though the rewards in both types of systems will not be available until the conversation is over, NL tutoring systems are more complex than the database-access dialogue systems described above. In dialogue systems like the train scheduler, the interaction time is much shorter, often less than 20 minutes, and the number of interactions within user-dialogue systems is generally less than 20 turns [Singh et al., 1999, Singh et al., 2002]. In NL tutoring systems, on the other hand, the preparatory training materials and testing typically exceed these timeframes significantly. In the studies reported here, it took students roughly 4-9 hours using the tutoring system itself, with around 280 interactions between a human subject and the NL tutoring system during the entire training sequence.

Additionally, compared with non-tutoring dialogue systems, there are two major challenges in applying RL to NL tutoring systems. Each of these will be discussed in the following section.

1.1.2.3 Whether RL Is Able To Induce Effective Tutorial Tactics Is Still An Open Question. The first main challenge is that it is difficult to determine which features of the learning environment are relevant, and thus, should be included in the state representation. Ideally the state should include all of the relevant dialogue history necessary to determine which action is taken next. One obvious but impractical choice is to use a complete record of the dialogue to the present point; however, in practice we need to compress the dialogue history to make the space tractable. The challenge lies in identifying the useful

features. Increasing the size of the state space may make the learning problem intractable, while the alternative may make the available data a much sparser sampling of the domain. While most of the work on the use of MDPs to improve dialogues has focused on developing the best policy given a set of features [Walker, 2000, Henderson et al., 2005], there has been relatively little work done on feature selection.

Early work on RL and MDP in non-tutoring dialogue systems focused largely on relatively simple task domains and used slot-based state representations. NJFun, for example, is a real-time spoken dialogue system that provides users with information about things to do in New Jersey. In applying RL to improve NJFun, Singh et al [Singh et al., 1999] used seven features for the state representation, such as whether the system has greeted the user, how many times a given piece of information has been asked for, and so on. However, as RL and MDP have been applied to more complex domains [Frampton and Lemon, 2006], the state space representations have increased in size and complexity, which creates a danger of making the learning problem intractable or the decision space too large to sample effectively.

Some of the previous studies in this area have focused on domain-specific features that should be included in the state-space. Singh et al. [Singh et al., 1999] showed that dialogue length was useful, while Frampton and colleagues [Frampton and Lemon, 2005, Frampton and Lemon, 2006] showed that incrementally adding high-level contextual information (such as the user’s last dialogue act and the last system move) into a state model, was also beneficial for building a better dialogue manager.

Previous research on applying RL in non-tutoring dialogue systems also investigated an effective feature selection procedure. Paek and Chickering’s work, for example, showed how a state-space can be reduced by only selecting features that are parents of the local immediate reward performs just as well as a more complicated model with other variables that are not parents [Paek and Chickering, 2005]. Recently, [Rieser and Lemon, 2006] used logistic regression to select the best state features for a multi-modal dialogue system and showed marked improvement over the baseline and some supervised learning methods. Most recently, Tetreault, et al [Tetreault and Litman, 2008] tackled the feature selection issue by exploring three evaluation metrics for assessing the utility of adding a particular state feature to a model of user state. The feature selection procedure employed in this dissertation is based

upon work by them. [Tetreault and Litman, 2008]. However, this thesis is fundamentally different than Tetreault et al’s work because they explored three evaluation metrics and used a relatively simple feature selection procedure. This thesis explored several different feature selection procedures, but used only one evaluation metric, the Expected Cumulative Reward (ECR). Specifically, Study 2 explored four categories of features suggested by the previous learning literature and a “greedy-like” feature selection method. Study 3 explored six categories of features and eleven more domain-general feature selection approaches.

The second main challenge is obtaining a training corpus. In order to use RL to induce an effective policy, it is necessary to collect an exploratory dataset that explores the relevant space of possible decision sequences. A common problem in RL is finding a balance between exploration (attempting to discover more about the world) and exploitation (using what we already know about the world to get the best results we can). A tutor in the real world must often choose between maximizing its expected utility according to its current knowledge about the world and trying to learn more about the world, since the latter may improve its future gains. This problem is known as the trade-off between exploitation and exploration.

Balancing exploration and exploitation is particularly important in educational contexts as data collection is generally very expensive. On one hand, without exploration, the tutor might not find an effective policy at all. On the other hand, if the tutor explores too much, it cannot stick to a path; in fact, it is not really learning as it cannot exploit its knowledge, and so acts as though it knows nothing. Thus, it is important to find a good balance between the two, to ensure that the tutor is really learning to take effective actions. It is often unclear how much exploration should be done in order to induce an effective policy. Ideally, of course, the training dataset should be as large as possible. One way to speed the process would be to use simulated student data [Beck, 2001, Beck et al., 2000, Ai and Litman, 2009]. Accurate simulations, however are difficult because the requirements for and causes of students’ learning are still open questions. An alternative approach is to use pre-existing data that was collected for other purposes. This route, however, is complicated by the fact that pre-existing systems often explore a small space and number of the actions and thus may yield biased or limited information.

In this dissertation, a different approach was taken. Instead of collecting a large ex-

ploratory training corpus at once, the corpus was accumulated over several stages. An initial exploratory dataset was collected that was large enough to apply RL to derive some tutorial policies. Once derived, those policies were used to train a new group of students and collect a new dataset. The new dataset was added to the original exploratory data to derive additional policies since the new group of students experienced the identical procedure as the original exploratory group. Such a process can be repeated until the learned policies either become stable or reach the desired results. It will be argued that this incremental improvement is the only practical method for continued improvement of an ITS. One would not want to continue using a poor quality tutoring system semester after semester when a better one could be available at the end of each semester. In this dissertation the process was repeated twice to determine if it resulted in an improved ITS. All data was collected before using the RL to adjust the tutoring system. In all, this dissertation includes three studies.

In Study 1, an initial NL dialogue system, called Cordillera, was built, in which the tutorial decisions on ET (elicit/tell) and JS (justify/skip-justify) were randomly made. This was used to collect an exploratory corpus by training a set of real students using the system. In Study 2, RL was used on the exploratory corpus to derive tutorial tactics, incorporate them back into Cordillera, train another group of students on the new version of system, and collect a new corpus. In Study 3, RL was applied to both the exploratory corpus from Study 1 and the new corpus from Study 2, individually and again on a single merged dataset combining students' corpus from both studies.

1.2 GENERAL APPROACH

As described above, there are two primary research questions in this thesis:

Question 1: Given content controlled among conditions, will micro-level pedagogical tutorial decisions affect student learning?

Question 2: Is RL a feasible method to induce tutorial tactics?

Each question represents a potential contribution to a field of research. Question 1 is relevant to the fields of learning and cognitive science; Question 2 is relevant to the fields of Intelligent Tutoring Systems, AI in Education, and Educational Data mining.

In order to investigate these two questions, Cordillera, a NL tutoring system was built which teaches students introduction to physics. Since Fall, 2007, three studies have been run [Jordan et al., 2007]. All three studies followed the same procedure: completing a background survey, reading a textbook, taking a pre-test, training on Cordillera, and finally, taking a post-test. All three studies used the same training problems and instructional materials but on different versions of Cordillera. The versions differed only in terms of the pedagogical tutorial tactics employed for micro-step level interactive decisions.

In Study 1, the Cordillera made interactive decisions randomly. This allowed us to collect an *exploratory corpus* that examined the consequences of each tutorial decision with real students. The student group for this study is referred to as the Exploratory Group. In order to differentiate this version of Cordillera from the ones used in subsequent studies, this version is referred to as Random-Cordillera.

In Study 2, RL was applied to the Exploratory corpus to induce a set of tutorial tactics named Dichotic Gain (DichGain) tutorial tactics. This version of Cordillera was named DichGain-Cordillera. DichGain-Cordillera employed the new policies to guide its interactive decisions. As before, this version of the system was used to train students in a complete study. The resulting corpus was named the *DichGain corpus* and the student group was named the *the Dichotic Gain (DichGain) group*. A preliminary analysis of these tactics, presented in Chapter 5, showed that they were no more effective than simple random decisions. In Study 3, RL was applied to induce tutorial tactics from both the Exploratory and DichGain corpora, both individually and again as a merged set. Two sets of tutoring tactics were derived from the three corpora, *Normalized Gain (NormGain)* and *Inverse Normalized Gain (InvNormGain)*. The NormGain set was derived with the goal of enhancing the tutorial decisions that contribute to the students' learning; while the InvNormGain set was derived with the goal of enhancing those decisions that contribute less or even none to the students' learning. We then ran a comparison study using the same educational materials as those in Studies 1 and 2. In Study 3 students were randomly assigned to one of two con-

ditions. One condition, the NormGain condition, was assigned to use a version of Cordillera which implemented the NormGain policies, named NormGain-Cordillera while another condition, the InvNormGain condition, was assigned to a another version of Cordillera with the InvNormGain policies, named InvNormGain-Cordillera.

Our primary hypothesis is:

**The Normalized Gain (NormGain) group will out-perform
the Inverse Normalized Gain (InvNormGain) group.**

The following thesis chapters will expand upon processes and outcomes from the studies summarized in Chapter 1. Chapter 2 provides a more detailed description of Cordillera, the Natural Language Tutoring System used in this thesis. Chapter 3 presents the detailed methodology for using the Reinforcement Learning toolkit to induce the dialogue management policies in this thesis. Chapters 4 through 7 present the three empirical studies. Chapter 4 focuses on collecting the Exploratory corpus. Study 2, collecting the DichGain corpus, is described in Chapter 5. Chapter 6 discusses the process for deriving the NormGain and the InvNormGain tutorial tactics for Study 3. Chapter 7 presents an experimental comparison of the induced NormGain and InvNormGain tutorial tactics in Study 3. Chapter 8 presents a general comparison across three studies and summarizes the conclusions. Chapter 9 discusses contributions to the fields of Cognitive and Learning Science, and to the fields of Artificial Intelligence and Education, Intelligent Tutoring Systems, and Educational Data Mining. Finally, this chapter considers future research initiatives that may evolve from this work.

2.0 CORDILLERA

This dissertation made use of the Cordillera system [VanLehn et al., 2007b]. Cordillera is a Natural Language (NL) based Tutoring System for introductory physics. The word “cordillera” is defined as an extensive range of mountains along a coastline, often consisting of a number of parallel chains. The Andes mountain range in South America is an example which includes the Cordillera Oriental and the Cordillera Occidental. As noted in Chapter 1, four different versions of the system were constructed, each of which differed only in terms of the tutoring tactics employed. Random-Cordillera, used in Study 1, made elicit/tell (ET) and justify/skip-justify (JS) decisions randomly; Dichotic Gain (DichGain) Cordillera, used in Study 2, followed DichGain policies induced from the exploratory corpus; while the Normalized Gain (NormGain) and Inverse Normalized Gain (InvNormGain) Cordillera systems, used in Study 3, followed the NormGain and InvNormGain policies induced from the Exploratory and DichGain corpora, individually or combined. The remaining components of the system, including the GUI interface, were identical for all participants.

Cordillera is based upon the TuTalk NL tutorial dialogue toolkit [Jordan et al., 2006, Jordan et al., 2007]. TuTalk is an authoring tool which enables domain experts to construct natural language tutoring systems without programming. Instead, the domain experts focus on defining the tutoring content through scripts, which are then used for automating interaction. TuTalk supports dialogues in which the tutor tries to elicit a line of reasoning from a student via a series of questions. This style of dialogue was inspired by the CIRCSIM-Tutor’s directed lines of reasoning [Evens and Michael, 2006]. In addition, TuTalk is modular, so that core modules, such as NL understanding, can be replaced or supplemented as needed.

To reduce confounds due to imperfect NL understanding in our experiments, the NL understanding module was replaced with a human interpreter called the language understanding

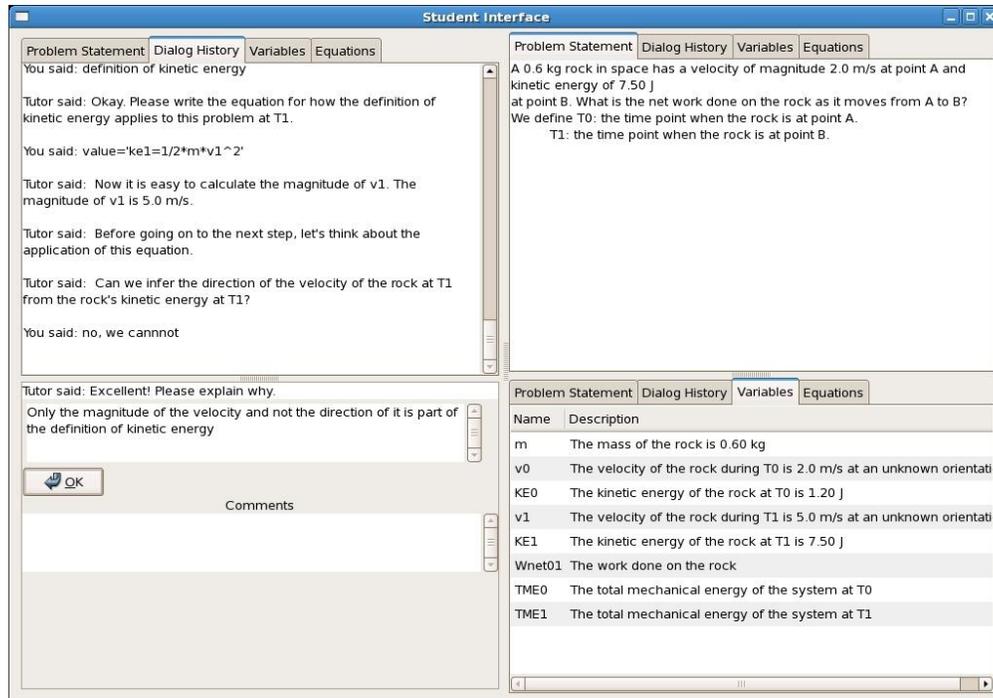


Figure 2.1: Student Interface

wizard [Bernsen and Dybkjaer, 1997]. In this format, Cordillera works as a communications framework that connects a student interface to a wizard interface. The student interface is used by students to read the tutor's tutorial instructions, to answer his or her questions, and to respond to them by means of natural language entries. The wizard interface is used to match students' answers to a list of potential responses. These two interfaces are discussed in detail below.

2.1 STUDENT INTERFACE

Figure 2.1 shows a screen shot of the student interface. The Message Window, located in the bottom-left corner is where the dialogue interaction takes place. The remaining four panes

are the Dialogue History Pane (upper-left), Problem Statement pane (upper-right), Variable Pane (lower-right) and the Equation Pane (not shown). The tabs included on three of the panels allow the student to select which four panels are visible and how where they will be displayed. Brief descriptions of each pane follow.

The *Message Window* is the focus of interaction between the student and tutor. All tutor messages appear here. Messages are displayed in this window, and students are then able to enter a response below. In some cases the response is merely an acknowledgement, i.e. clicking the [OK] button to proceed to the next action. In other cases, such as when the tutor asks a question, the student can submit an answer by typing in a text field.

The *Dialogue History Pane* shows a record of the student-tutor dialogue thus far.

The *Problem Statement Pane* shows the problem statement and any accompanying figures if present.

The *Variable Description Pane* shows all the variables defined during problem solving. The variables can be defined either by the student using a form interface or provided by the tutor (elicit vs. tell).

Finally, the *Equation Pane* displays the equations that have been input either by students or by tutors up to that point in the problem solving. An equation is presented as a two-column table where each row consists of a formula and its description. A description consists of the name of the principle and its arguments.

2.2 WIZARD INTERFACE

The Wizard Interface, shown in Figure 2.2, mirrors the student interface in all respects with the exception of the Message Window. This is replaced by the Student Response Classification Window, which displays the student's most recent response along with a set of check-boxes for classifying the response. In the example below, the student's response was classified as the third choice. Once the student's response has classified, the system would follow control scripts to decide what to do next, and the dialogue manager would decide how to do it. If none of the choices match the student's entry, then the wizard makes no selection

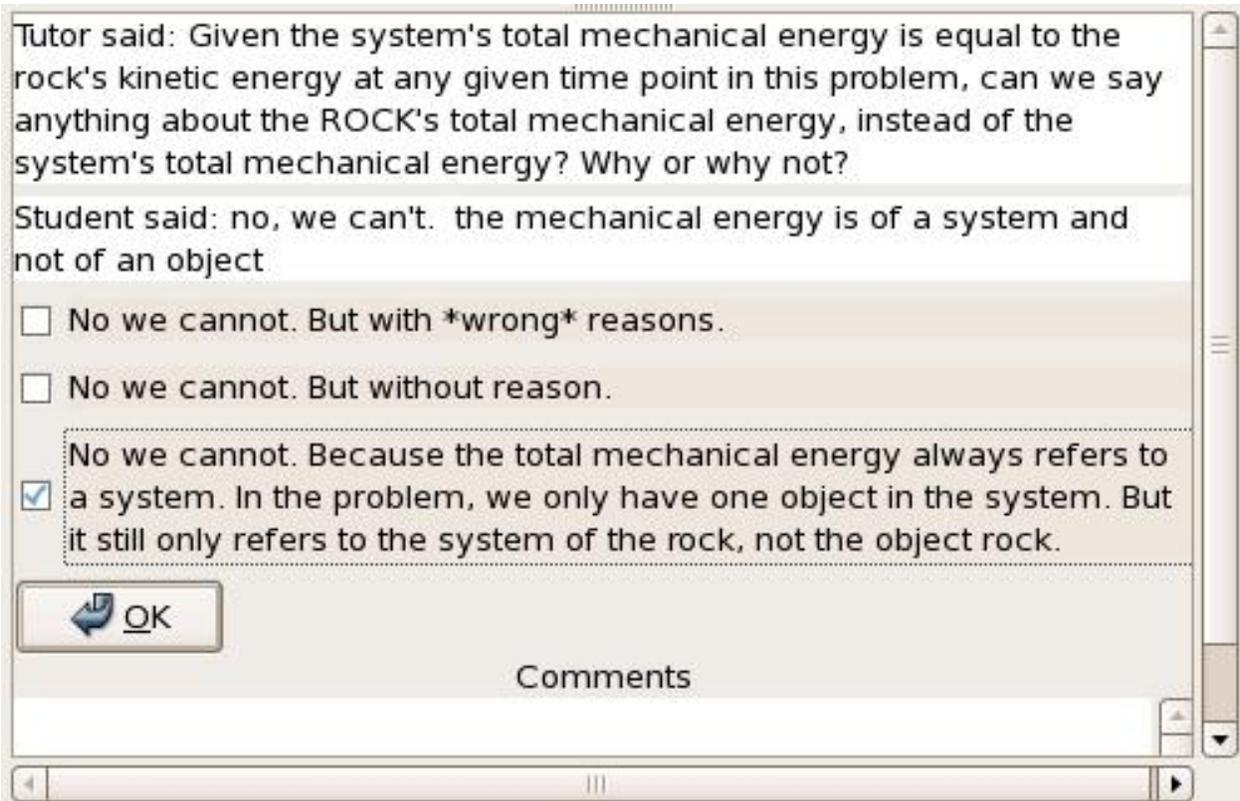


Figure 2.2: Students' Response Classification Window in Wizard Interface

and simply clicks the OK button.

2.3 AN EXAMPLE SCRIPT

Cordillera dialogues are governed by control scripts authored by domain experts. These scripts control the dialogue messages as well as the content of each information panel. An example script is shown in Table 2.1. In the example script, the highest level step is “define-system” in Line 1. By name, it is about defining a system. This step consists of three micro-steps listed in the order of Lines 2, 3, and 4.

Table 2.1: A Sample Cordillera Script

1.	<u>g define-system</u>
2.	do choose-system SEM ELICIT/TELL
3.	do system-justification SEM JUSTIFY; ELICIT/TELL
4.	do isolated-system SEM ELICIT/TELL
5.	<u>g choose-system SEM ELICIT</u>
6.	say “What would be your choice of the system for this problem?”
7.	<i>if</i> “truck-Earth.” true
8.	<i>if</i> “truck and the Earth.” true
9.	<i>otherwise</i> do bottom_out_choose-system
10.	<u>g bottom_out_choose-system</u>
11.	say “What are the object(s) in the problem?”
12.	<i>if</i> “truck and the Earth” true
13.	<i>otherwise</i> say “[cont]There are two objects here, the truck and the Earth.”
14.	say “[cont]The best choice of the system here is to select both.”
15.	<u>g choose-system SEM TELL</u>
16.	say “[cont]There are two objects in the problem, the truck and the Earth. The best choice of the system for this problem is to choose both. “

Line 2 shows that the tutor needs to make an Elicit/Tell (ET) decision for the micro-step “choose-system” since it is labeled with “SEM ELICIT/TELL.” Here “SEM” stands for the term “semantic”. This is a feature of Tutalk, which was meant to be used to mark semantically similar turns and to allow the student model to make decisions relative to that semantically similar content. Line 3 shows that the tutor needs to make two decisions on the

micro-step “system-justification” since it is labelled with “SEM JUSTIFY; ELICIT/TELL”. When a tutorial decision step involves both ET and JS decisions, the system always makes the JS decision first. If it decides to skip the justification step, then the system does not need to make the ET decision and goes to next micro-step “isolated-system” in Line 4. On the other hand, if it decides not to skip the “system-justification”, the system would then make the ET decision. Line 4 shows that the tutor needs to make the ET decision on the micro-step “isolated-system”. The next paragraph describes how the elicit and tell versions of the “choose-system ” are executed in Cordillera.

Executing the *Elicit* Version of “choose-system”: If the system decides to elicit on the micro-step “choose-system ” in Line 2, then the scripts from Lines 5-14 will be executed, because they are the elicit version of “choose-system. ” Lines beginning with the command “say”, that is, lines 6, 11, 14, and 16, initiate a tutor message. For each “say” line, if the content immediately following it is a normal sentence, as in lines 6 and 11, then the tutor will send the text to the student as a question, requiring the student to respond. For example, when TuTalk sends Line 6 to Cordillera, it will then display the question “What would be your choice of the system for this problem?” in the Message Window on the student’s interface along with a text field for the student’s answer. The question will also be shown in the Student Response Classification Window on the Wizard Interface. Once the student inputs an answer and clicks OK, that answer will be added to the Student Response Classification Window in the Wizard Interface along with the set of possible choices as a checklist. These choices are taken from the lines prefixed with *if* below the “say.” In the case shown in Figure 2.1 the answer choices for the “say” in Line 6 are “truck-earth” in Line 7 or “truck and the earth” in Line 8. The human wizard chooses from this list the answer that is closest to the student’s answer and submits it. Unmatched answers are left blank. Both the student’s answer and the wizard’s match are sent to the system to guide the next decision.

Executing the *Tell* Version of “choose-system”: On the other hand, if the system decided to tell on “choose-system ”, then the tell version of “choose-system” in Lines 15-16 will be executed. If the contents after “say” is a sentence starting with [*cont*], which rep-

resents “continue,” as in lines 14 and 16, then the tutor will tell the sentence following the [cont]. For example, when line 16 is sent to Cordillera by TuTalk, the Message Window on the student’s Interface will show the tutor’s message “There are two objects in the problem, the truck and the Earth. The best choice of the system for this problem is to choose both.” followed by an [OK] button. The student acknowledges the dialogue by clicking the OK button at which point the dialogue will move on to the next topic.

Executing Justify/Skip-justify: For Justify/Skip-justify (JS) decisions, the scripts are much simpler. An example of this is listed in Line 3. The label “JUSTIFY” in 3 indicates that this is a JS tutorial decision step. If the dialogue manager decides to execute the micro-step *system-justification* in Line3, the system will make the next ET decision on the step; otherwise, the dialogue manager skips this micro-step and goes directly to execute the next micro-step “isolated-system” in Line 4.

To summarize, the design of Cordillera allows domain experts to manage a natural-language tutorial dialogue including GUI components such as variable listings and an equation display. The scripts allow variables and equations to be added or removed as needed as well as for selected portions of the display to be highlighted. Alternative dialogue actions such as elicits and tells, justify and skip-justify, are encoded in the dialogue for selection by the dialogue manager. The script authors determine the flow of the dialogue within these alternatives and the content of each question, including alternative choices. The NL tutoring system or the wizards in the thesis study match these alternatives to student input at runtime.

At present, the dialogue scripts can be written in XML or other more readable formats, which are then compiled by the TuTalk system into executable script form or other textual formats. TuTalk will be able to translate these human readable files into TuTalk- and Cordillera-readable ones. Moreover, the reverse is also possible. As a result, the domain experts can build an NL tutoring system by simply focusing only on the subject matter.

This chapter illustrated the Cordillera system. Chapter 3 will present a detailed description of the research methodology focusing on how to apply reinforcement learning to induce the tutorial policies in this thesis.

3.0 REINFORCEMENT LEARNING PROCEDURE

Chapter 3 further develops the previous discussion of RL and describes the general procedure by which tutorial dialogue policies were induced from student interactivity data. The chapter begins with a description of how a problem of inducing pedagogical tutorial tactics can be fit into the general RL and MDP framework. In this thesis it is assumed that inducing tutorial tactics specific to each Knowledge Component (KC) will be more effective than inducing an overall KC-general policy. This chapter provides an overview of KCs (a more detailed description of the identified KCs may be found in Chapter 4) and the approach used to generate KC-based MDPs from the training corpus. It also describes the induction toolkit employed and the assessment metrics used. Finally, the chapter discusses the issues confronted when the induced policies were implemented back into Cordillera.

In this dissertation, a toolkit is used to calculate an optimal dialogue policy given a suitable Markov Decision Process (MDP) model. The major challenge faced therefore was the production of the MDP model, especially the KC-based MDPs. There were five distinct issues that needed to be addressed. The list included collecting and/or selecting training corpora from which the tutorial tactics will be derived, determining on which KCs the tutorial tactics should be induced for, the reward function, the state representation, and how to handle conflicting policies. These issues are discussed below. Additional details about how these issues are addressed in Studies 2 and 3 and the resulting models are presented in Chapters 5 and 6.

3.1 REINFORCEMENT LEARNING FOR TUTORIAL TACTICS

Previous research on using RL to improve dialogue systems (e.g. [Levin and Pieraccini, 1997, Singh et al., 1999]) has typically used MDP’s [Sutton and Barto, 1998] to model dialogue data. The central idea behind this approach is to transform the problem of inducing effective dialogue policies into computing an optimal policy for choosing actions in an MDP. An MDP formally corresponds to a 4-tuple (S, A, T, R) , in which:

$S = \{S_1, \dots, S_n\}$ is a state space.

$A = \{A_1, \dots, A_m\}$ is an action space represented by a set of action variables;

$T : S \times A \times S \rightarrow [0, 1]$ is a set of transition probabilities between states that describe the dynamics of the modeled system; for example: $P(S_j|S_i, A_k)$ is the probability that the model would transition from state S_i to state S_j by taking action A_k .

$R : S \times A \times S \rightarrow R$ denotes a reward model that assigns rewards to state transitions and models payoffs associated with such transitions.

Additionally, $\pi : S \rightarrow A$ is defined as a policy or tutorial tactics.

Dialogue management can be easily represented using an MDP: the states are vector representations composed of relevant student and dialogue characteristics; the transitions are dialogue system acts; and the reward function is calculated from the dialogue system’s success measures such as completion on task, and, in the present case, learning gains. More formally, we can view each dialogue d_i as a trajectory in the chosen state space determined by the system actions and user responses:

$$s_{d_i}^1 \xrightarrow{a_{d_i}^1, r_{d_i}^1} s_{d_i}^2 \xrightarrow{a_{d_i}^2, r_{d_i}^2} \dots s_{d_i}^{n_{d_i}} \xrightarrow{a_{d_i}^{n_{d_i}}, r_{d_i}^{n_{d_i}}}$$

Here $s_{d_i}^j \xrightarrow{a_{d_i}^j, r_{d_i}^j} s_{d_i}^{j+1}$ indicated that at the j th turn in the dialogue d_i , the system was in state $s_{d_i}^j$, executed action $a_{d_i}^j$, received reward $r_{d_i}^j$, and then transferred into state $s_{d_i}^{j+1}$. The number of turns in d_i is n_{d_i} . For a training corpus consisting of L dialogues with n_1, n_2, \dots, n_L turns in each dialogue respectively, the training corpus looks like:

$$\begin{aligned} & s_{d_1}^1 \xrightarrow{a_{d_1}^1, r_{d_1}^1} s_{d_1}^2 \xrightarrow{a_{d_1}^2, r_{d_1}^2} \dots s_{d_1}^{n_1} \xrightarrow{a_{d_1}^{n_1}, r_{d_1}^{n_1}} \\ & \dots \\ & s_{d_i}^1 \xrightarrow{a_{d_i}^1, r_{d_i}^1} s_{d_i}^2 \xrightarrow{a_{d_i}^2, r_{d_i}^2} \dots s_{d_i}^{n_i} \xrightarrow{a_{d_i}^{n_i}, r_{d_i}^{n_i}} \\ & \dots \\ & s_{d_L}^1 \xrightarrow{a_{d_L}^1, r_{d_L}^1} s_{d_L}^2 \xrightarrow{a_{d_L}^2, r_{d_L}^2} \dots s_{d_L}^{n_L} \xrightarrow{a_{d_L}^{n_L}, r_{d_L}^{n_L}} \end{aligned}$$

Once the MDP structure S, A, R has been defined, the model parameters T are estimated from the training corpus as: $T = \{p(S_j|S_i, A_k)\}_{i,j=1,\dots,n}^{k=1,\dots,m}$. More specifically, it is calculated by taking the number of times that the system is in state S_i , took step A_k , and arrived in state S_j divided by the number of times the system was in S_i and took A_k . The reliability of these estimates clearly depends upon the size and structure of the training dataset. Once a complete MDP is constructed, a dynamic programming approach can be used to learn the optimal control policy π^* , i.e. the set of actions the model should take at each state, to maximize its expected cumulative reward.

3.2 ISSUE 1: TRAINING CORPUS

One of the main characteristics that differentiate RL from other machine learning techniques is exploration. In order to have confidence in the constructed MDP, the training corpus must explore various possible actions from various possible states, and preferably, many times. In other words, the training corpus must be exploratory with respect to the chosen states and actions. If we never try an allowed action from a certain state, we cannot expect to know the value of taking that action in that state. As a result, unexplored state transitions cannot be estimated, and transitions that are explored infrequently will have poor or strongly biased estimates. Even a large, but biased, corpus presents problems in that it may focus extensively on one small subset of the domain. This is especially true with pre-existing tutoring corpora where the data is gathered by using a system with a hand-tooled rule set. In that situation, the existing policy and the subjects' use of it may bias the dataset and prevent adequate exploration. Some authors have proposed using simulated students to generate training data [Levin and Pieraccini, 1997, Young, 1999]. It is still an open question, however, about what causes students to learn and how they learn. As a result, constructing a valid simulation that provides an accurate estimate of students' responses and their learning is doubtful. Therefore this thesis focused solely on real user data. Singh et al. suggest that authors avoid using biased data by collecting "exploratory data," that is, data collected from a system that makes tutorial decisions randomly, thus ensuring that the transitions are adequately

explored [Singh et al., 2002].

In this research two approaches were adopted. In Study 1 students made use of Random-Cordillera which made the crucial Elicit/Tell and Justify/Skip-justify decisions randomly. In Study 2 students made use of Dichotic Gain (DichGain) Cordillera which made decisions based upon a policy induced from the training corpus collected in Study 1. The former random route is consistent with the exploration literature, while the latter route is consistent with the task of gradually improving induced policies over time. As will be described in Chapter 6, in preparation for Study 3 both corpora were used for policy induction.

3.3 ISSUE 2: KNOWLEDGE COMPONENTS

In tutoring literature, it is commonly assumed that relevant knowledge in domains such as math and science is structured as a set of independent but co-occurring Knowledge Components (KCs). A KC is “a generalization of everyday terms like concept, principle, fact, or skill, and cognitive science terms like schema, production rule, misconception, or facet” [VanLehn et al., 2007b]. For the purposes of tutoring systems, these are the atomic units of knowledge. Problem solving in such domains typically involves complex problems consisting of multiple steps, each of which involves a single or a combination of independent KCs. For example, a simple algebraic equation, $2x + 5 = 21$ can be solved via two steps: 1) subtract the same term 5 from both sides of the equation; and 2) divide both sides by the non-zero term 2. Here subtracting the same term from both sides of the equation is one KC and dividing both sides of the equation by the non-zero term is another KC. As problems grow more complex, the number of KC’s involved, and their combinations, can increase exponentially.

In the tutoring literature it is commonly assumed that KC’s are learned independently of one another. A number of standardized tests, for example, are constructed based on this assumed independence among KCs. Techniques exist to re-engineer the definition of KCs so that they are independently learnable [Cen et al., 2006, Cen et al., 2007], thus improving the overall effectiveness of the resulting tutoring system. When dealing with a specific KC, the expectation is that the tutor’s interactive decision on that KC, elicit or tell, would

be based upon the student’s mastery of the KC in question, its intrinsic difficulty, and other relevant, but not necessarily known, factors specific to that KC. In other words, the assumption is that an optimal policy for one KC might not be optimal for another. Therefore, the assumption made in this dissertation is that inducing tutorial tactics specific to each Knowledge Component (KC) would be more effective than inducing an overall KC-general policy. The KCs identified by the domain experts for the domain are described below. In order to derive KC-based tutorial tactics, KC-based MDPs needed to be generated from the training corpus.

3.3.1 Identified KCs in the Selected Domain

In order to learn a policy for each KC, the KCs in a domain need to be identified. The domain chosen for this dissertation covers the work-energy chapter in college-level physics textbook. Two domain experts (not the author) who are also knowledge representation experts identified 32 KCs in the domain. For example, KC_{20} and KC_{21} are the two KCs that were involved in the majority of the tutorial decisions on elicit/tell (ET) and justify/skip-justify (JS) respectively.

Definition of Kinetic Energy ($KE = \frac{1}{2}mv^2$) — KC_{20} : If an object is moving, then its kinetic energy at a time is $\frac{1}{2}mv^2$, where m is the object’s mass and v is the magnitude of the object’s instantaneous velocity.

Definition of Gravitational Potential Energy ($GPE = mgh$) — KC_{21} : If an object and a planet are in a system (or equivalently, the gravitational force of the earth on the object is an internal force), then their gravitational potential energy is mgh , where m is the mass of the object, g is the gravitational acceleration of the planet, and h is the object’s height above a zero point. The zero point is arbitrary, but is often chosen to be the planet’s surface.

Note that a complicated domain like physics can often be broken into many KCs. Here the 32 identified KCs are believed to cover the most important knowledge in the domain. There are some other KCs shown in the tutorial decision steps that are not among 32 identified KCs. After identifying the KCs involved in the domain, we needed to decide which KCs are

needed in order to induce KC-based policies. Intuitively, it should depend on the relative importance of these KCs and also on the frequency of their appearances in the training corpus, and so on. To determine which KCs were required to induce the KC-based tutorial tactics, the training tutorial dialogues were annotated with the tutorial action decisions based on the KCs involved.

3.3.2 Tutorial Dialogue Annotation

A group of five individuals (including the author) annotated each of the tutoring dialogues and action decisions with the relevant KCs. The KCs were drawn from the set of 32 KCs described in Appendix A. Each tutorial dialogue consists of one human participant’s interaction with Cordillera to solve 7 problems. For each of seven problems, there were at least two annotators. For each of 32 identified KCs, the final kappa was ≥ 0.77 which is fairly high given the complexity of the task.

Table 3.1 shows an example. Here lines 1, 5, 6, 7, 9 and 11 are tutor turns with the action decisions and their target KCs listed on the right. Next, various labels in Table 3.1 will be explained.

Line 1 was labelled “(**ELICIT**, KC_{22})”. This indicates that this is an ET decision step on KC_{22} and the tutor decided to elicit.

Line 2 was labelled “(+ KC_{22})”. It shows that the student answered the question correctly, so a + KC_{22} was added to the student’s turn. Or it is a correct entry generated by the student. Alternatively, a $-KC_{22}$ would represent that it is an incorrect entry input by the student on KC_{22} .

Line 5 was labelled “(**TELL**, $*KC_{22}$)”. It indicates that the tutor made an ET decision on KC_{22} and decided to tell the student directly. Here “ $*KC_{22}$ ” indicates that this piece of information is about KC_{22} , and it was generated by the tutor. By contrast, ones that were generated by the students are labelled with “+” and “-” to represent correct and incorrect entries respectively.

Line 6 was labelled with a simple “(**TELL**)”. It indicated that the tutor made an ET decision, but this step *does not involve any identified KCs*.

Table 3.1: Example of Tutorial Dialogue with KC Labels

1. *T*: At T2 we also know the spring is uncompressed. So, what type of energy can we easily find now? (**ELICIT**, KC_{22})
2. *S*: spe (+ KC_{22})
3. *T*: Yes, I agree.
4. *T*: At this point we know the values for KE2 and GPE2. It is time to compute the value for SPE2.
5. *T*: It is important to remember that we cannot find the SPE of an OBJECT. We need to find the spring potential energy of a pair: object-spring. (**TELL**, $*KC_{22}$)
6. *T*: It is possible to compute the spring potential energy at T2 because we know the spring is uncompressed. That is, $d_2=0$ m. (**TELL**)
7. *T*: Using this information we can infer that the spring potential energy of the block-spring pair at T2 is $SPE_2=0$ J. (**TELL**, $\{*KC_{22} *KC_{23}\}$)
8. *T*: Now we have all we need to find the value for the system's total mechanical energy at T2.
9. *T*: TME of the system at T2 is the sum of the block's KE and the block-Earth pair's GPE because the value for the spring potential energy is zero. In this case $TME_2=11.88$ J+1.62 J =13.5 J. (**TELL**, $\{*KC_{23} *KC_{24}\}$)
10. *T*: At this point we have found the value for TME2. We can now easily find the system's TME at all other time points since the system is isolated.
11. *T*: What principle gives us the system's total mechanical energy at every time point, knowing its value at T2 and also given that the system is isolated? (**ELICIT**, $\{KC_{27}\}$)
12. *S*: student conservation of energy (+ KC_{27})

Line 9 was labelled with “(**TELL**, $\{*KC_{23} *KC_{24}\}$)”. It indicated that the tutor made an ET decision and that this step involves two KCs: $\{KC_{23} KC_{24}\}$). Because the tutor told the information, both KCs are labeled with “*”.

The last tutor turn in the example is Line 11, labelled “(ELICIT, {KC₂₇})”. This indicates that the tutor made an ET decision on KC₂₇ and decided to elicit the information from the student.

Finally, Line 12 was labelled “(+KC₂₇)”. It indicated that it was a student turn and the student input a correct entry for KC₂₇.

For a training dialogue d_i , one dialogue trajectory can be constructed for each KC. More specifically, we use n_{d_i, KC_k} to represent the number of turns on KC_k in the dialogue d_i and we expect different n_{d_i, KC_k} for different KC_k because the number of tutorial actions on each KC varies in the d_i . Thus, we have:

$$\begin{aligned}
& s_{d_i, KC_1}^1 \xrightarrow{a_{d_i, KC_1}^1, r_{d_i, KC_1}^1} s_{d_i, KC_1}^2 \cdots s_{d_i, KC_1}^{n_{d_i, KC_1}} \xrightarrow{a_{d_i, KC_1}^{n_{d_i, KC_1}}, r_{d_i, KC_1}^{n_{d_i, KC_1}}} \\
& \quad \quad \quad \cdots \\
& s_{d_i, KC_k}^1 \xrightarrow{a_{d_i, KC_k}^1, r_{d_i, KC_k}^1} s_{d_i, KC_k}^2 \cdots s_{d_i, KC_k}^{n_{d_i, KC_k}} \xrightarrow{a_{d_i, KC_k}^{n_{d_i, KC_k}}, r_{d_i, KC_k}^{n_{d_i, KC_k}}} \\
& \quad \quad \quad \cdots \\
& s_{d_i, KC_P}^1 \xrightarrow{a_{d_i, KC_P}^1, r_{d_i, KC_P}^1} s_{d_i, KC_P}^2 \cdots s_{d_i, KC_P}^{n_{d_i, KC_P}} \xrightarrow{a_{d_i, KC_P}^{n_{d_i, KC_P}}, r_{d_i, KC_P}^{n_{d_i, KC_P}}}
\end{aligned}$$

Here P is the number of the KCs used to induce policies, and they are represented as: KC_1, \dots, KC_P . And $s_{d_i, KC_k}^j \xrightarrow{a_{d_i, KC_k}^j, r_{d_i, KC_k}^j} s_{d_i, KC_k}^{j+1}$ indicated that in the j th turn in the dialogue d_i for KC_k , the system was in state s_{d_i, KC_k}^j , executed action a_{d_i, KC_k}^j , received reward r_{d_i, KC_k}^j , and then transferred into state s_{d_i, KC_k}^{j+1} .

For each KC_k , all of the dialogue trajectories on KC_k from each tutorial dialogue were combined in the training corpus and used to generate the training corpus for KC_k . From that training corpus, an MDP_{KC_k} could be constructed and a dynamic programming approach could be used to learn the optimal control policy $\pi_{KC_k}^*$. In the next section, the procedure is described in detail.

3.4 KC-BASED MDPS

In order to induce KC-based tutorial tactics, an MDP model was constructed for each KC. For example, for KC KC_k , the corresponding MDP_{KC_k} is defined as:

$S_{KC_k} = \{S_{KC_k,1}, \dots, S_{KC_k,n}\}$ correspond to the dialogue states related with KC_k
 $A_{KC_k} = \{A_{KC_k,1}, \dots, A_{KC_k,m}\}$ correspond to the tutorial actions involving KC_k only.
 $T_{KC_k} : S_{KC_k} \times A_{KC_k} \times S_{KC_k} \rightarrow [0, 1]$ is a set of transition probabilities between KC_k related states and action A_{KC_k} on KC_k . Once the MDP_{KC_k} structure has been defined, they are estimated from the corresponding annotated training corpus on KC_k .
 $R_{KC_k} : S \times A_{KC_k} \times S_{KC_k} \rightarrow R_{KC_k}$ are defined to reflect dialog performance metric on KC_k only.
Additionally, $\pi_{KC_k}^* : S_{KC_k} \rightarrow A_{KC_k}$ is defined as a KC-based policy or tutorial tactics for KC_k .

The general approach for defining KC-based state representation, action choice space, and reward function is described in the next subsection of this dissertation.

3.4.1 Issue 3: State Representation

For RL, as with all machine learning tasks, success is dependent upon choosing an appropriate set of features to represent dialogue states. An ideal state representation should include all of the tutorial dialogue information that is relevant and necessary to determine what action the system should take next. Ideally this would include a complete record of the tutoring interaction thus far—both for the present problem and preceding problems—as well as derived features such as gender, MSAT, detailed pre-test scores, and so on. However, the high cost of obtaining human tutorial dialogues makes it crucial to limit the size of the state space. Nevertheless, even a state based on a handful of features can yield an enormous state space.

In order to obtain an effective representation that both minimizes data sparsity while retaining sufficient information, a small but carefully selected feature space is preferable. Using a small state representation to approximate the true state reduces the amount of data required. The disadvantage of doing so is that it increases the risk that educationally relevant features will be missed, resulting in a non-representative state space. To this end this thesis began with a large set of features to which a series of feature-selection methods were applied

to reduce them to a tractable subset. Because of this, the state representation issue can be divided into four sub-issues for discussion purposes.

3.4.1.1 Sub-Issue 1: Feature Choices This sub-issue concerned what types of relevant information can be included in the state space. For this dissertation only features that could be computed automatically or evaluated objectively, such as gender, were included. Hand-annotated dialogue features were omitted as the tutor would require the features to be available in real time when the learned policies are employed. Moreover, in order to induce KC-based tutorial tactics, the state representations were also KC-based. For example, “pctCorrectKCPM” is a feature choice in Study 2. It is defined as the students’ performance on the specific KC. In $MDP(KC_{20})$ it refers to the students’ performance on KC_{20} , while in $MDP(KC_{21})$ it refers to the students’ performance on $MDP(KC_{21})$.

3.4.1.2 Sub-Issue 2: Feature Discretization An MDP model generally requires all the state features in the model to be discrete variables. Most of the features of interest here, such as “pctCorrectKCPM”, are continuous. It is thus necessary to choose an effective method for feature discretization.

3.4.1.3 Sub-Issue 3: Feature Selection One of the main challenges in this dissertation was feature selection. For Study 2, a greedy-like search strategy for feature selection was employed, while in Study 3 more extensive feature selection methods were employed. These will be discussed in more detail in Chapters 5 and 6 respectively.

3.4.1.4 Sub-Issue 4: Maximum Number of Features The last major sub-issue related to state representation is the maximum number of features to be included in the state space. The number should be small so that we have enough training data to cover each state, yet large enough to would include enough features to represent states without losing information necessary to make good tutorial decisions.

3.4.2 KC-based Action

In the present studies there are two types of choices in the action space: Elicit/Tell (ET) and Justify/Skip-justify (JS). These are available to the tutor at different times for different KCs in the tutoring process. In each $MDP(KC_k)$, the tutorial action choices were those that involved the specific KC_k only.

3.4.3 Issue 4: KC-based Reward

Based on previous research by [Tetreault and Litman, 2008, Tetreault and Litman, 2006b, Tetreault and Litman, 2006a] Normalized Learning Gain (NLG) was selected as a reward function because it measures students' gain irrespective of their incoming competence. In addition to mapping the training problems to KCs, a domain expert also mapped the pre/post test problems to the sets of relevant KCs. This resulted in a KC-specific NLG score for each student. The reward function is defined as follows:

$$NLG_{KC_k} = \frac{posttest_{KC_k} - pretest_{KC_k}}{1 - pretest_{KC_k}} \quad (3.1)$$

Here $post - test_{KC_k}$ and $pre - test_{KC_k}$ refer to the KC-specific pre- and post-test scores on KC_k for each student.

In this dissertation, only terminal dialogue states have non-zero rewards because a student's NLG will not be available until the entire of his/her tutorial dialogue is completed. Thus for a tutorial dialogue d_i , $r_{d_i, KC_k}^1 \dots, r_{d_i, KC_k}^{n_{d_i, KC_k} - 1}$ are all equal to 0 and only the final reward $r_{d_i, KC_k}^{n_{d_i, KC_k}}$ equals to non-zero rewards. The final reward in this thesis is determined by the student's NLG on the corresponding KC. Here n_{d_i, KC_k} represents the number of turns that tutorial dialogue d_i had made decisions on KC_k .

3.5 INDUCE KC-GENERAL POLICIES

To this point, the dissertation has focused on inducing KC-based tutorial tactics. However, certain tutorial decision steps do not involve any *identified* KCs. Line 6 in Table 3.1 is

such an example. Next we need to decide how the dialogue manager should perform in this instance. In this dissertation, the issue was resolved by inducing a KC-general tutorial tactic.

In both Study 2 and Study 3, we induced one KC-general policy for ET and one for JS. This was done by using the same general approach as inducing KC-based policies except that the state representation, action, and choice are no longer based on any particular KC. For KC-general policies, the final rewards are calculated based upon the student’s cumulative KC-based NLGs. When a tutorial decision step does not involve any KCs, the dialogue manager would follow the KC-general policies.

Once an MDP model has been completed, calculation of an optimal policy is straightforward. This dissertation work employed an RL toolkit developed by Tetreault and Litman [Tetreault and Litman, 2008, Tetreault and Litman, 2006b, Tetreault and Litman, 2006a].

3.6 TETREULT AND LITMAN’S RL TOOLKIT

Tetreault, & Litman’s toolkit [Tetreault and Litman, 2008, Tetreault and Litman, 2006b, Tetreault and Litman, 2006a] uses a dynamic programming algorithm for policy iteration [Sutton and Barto, 1998]. The code was originally built on the MDP toolkit written in Matlab [Chades et al., 2005]. The purpose of this algorithm is to handle the problem of reward propagation. As noted above, rewards, in this case learning gains, are not assigned until the end of the tutoring process, long after any action has occurred. The dynamic programming algorithm propagates the rewards back to the internal states weighting the V-value of each state, s via the following recursive equation:

$$V(s) = \max_a R(s, a) + \sum_{s'} P(s'|s, a)\gamma V(s') \quad (3.2)$$

Here $P(s'|s, a)$ is the estimated transition model from the training corpus, $R(s, a)$ is the estimated reward model, and $0 \leq \gamma \leq 1$ is a discount factor. If γ is less than 1, then it will discount rewards obtained later. For all the studies reported here, a discount factor of 0.9 was used, which is common in other RL models [Tetreault and Litman, 2008, Tetreault and Litman, 2006b, Tetreault and Litman, 2006a].

The V-values, as defined by Equation 3.2, can be estimated to within a desired threshold using policy iteration [Sutton and Barto, 1998]. Here an estimated v-value and a best possible action to take for each state are recorded. These are then iteratively updated based on the values of its neighboring states. This iteration stops when each update yields a difference below some threshold ϵ . Once the policy iteration process is complete, the optimal dialogue policy π^* is obtained by selecting the action that produces the highest expected reward (or V-value) for that state. At this time we also compute the Expected Cumulative Reward (ECR) and a 95% confidence interval for the ECR (hereafter, 95%CI) for the optimal policy.

3.6.1 Expected Cumulative Reward (ECR)

The Expected Cumulative Reward (ECR) of a policy is derived from a side calculation in the policy iteration algorithm: the V-values of each state, the expected reward of starting from that state and finishing at one of the final states. More specifically, the ECR of a policy π can be calculated as follows:

$$ECR_{\pi} = \sum_{i=1}^m \frac{N_i}{N_1 + \dots + N_m} \times V(s_i) \quad (3.3)$$

Where s_1, \dots, s_m is the set of all starting states and $v(s_i)$ is the V-values for state s_i ; N_i is the number of times that s_i appears as a start state in the model and it is normalized by dividing $\frac{N_i}{N_1 + \dots + N_m}$. In other words, the ECR of a policy π is calculated by summing over all the initial start states in the model space and weighting them by the frequency each state appears as a start state.

In Tetreault and Litman’s work [Tetreault and Litman, 2008], the authors used ECR as a evaluation metric for feature selection. Additionally, ECR has been widely used as the criteria for evaluating the policy in the area of inducing policy from simulated corpus [Janarthanam and Lemon, 2009, Williams and Young, 2007b, Williams and Young, 2007a]. More specifically, given two MDP structures: $MDP_1 = \{S_1, A, R\}$ and $MDP_2 = \{S_2, A, R\}$, which have the same action choices A and reward function R but different state representation, the transition probability T_1 and T_2 were estimated from the *same* training corpus. Two different policies, π_1 and π_2 , were derived based on MDP_1 and MDP_2 respectively.

The higher the ECR value of a policy, the better the policy is supposed to perform.

3.6.2 Confidence Interval

Tetreault and Litman pointed out one limitation of using the ECR as an evaluation metric for a policy: it assumes that there was sufficient collected data to derive a reliable policy [Tetreault and Litman, 2008, Tetreault and Litman, 2006b, Tetreault and Litman, 2006a].

However, in practice researchers frequently have to deal with issues of data sparsity. They proposed a novel approach by taking into account the reliability of the transition probability estimates from the training data and constructing a confidence interval for the ECR for the learned policy.

As described earlier, an estimate for the ECR was computed by using the transition probabilities derived from the training corpus. Note that these transition probabilities are simply estimates which are more or less accurate, depending on how much data is available. As an illustration, Tetreault and Litman used the following example [Tetreault and Litman, 2008]: In an MDP model, we have $S = \{S_1, S_2, S_3\}$, $A = \{A_1, A_2\}$. From a training corpus, there were ten cases that an action A_1 was taken from state S_1 . Out of these, three times the system transitioned back to state S_1 , two times it transitioned to state S_2 , and five times to state S_3 . Thus we have

$$P(S_1|S_1, A_1) = \frac{3}{10} = 0.3 \tag{3.4}$$

$$P(S_2|S_1, A_1) = \frac{2}{10} = 0.2 \tag{3.5}$$

$$P(S_3|S_1, A_1) = \frac{5}{10} = 0.5 \tag{3.6}$$

From the same corpus, there were 1000 times that action A_2 was taken from state S_2 . In 300 of those cases it transitioned to state S_1 ; in 200 cases to state S_2 ; and the remaining

500 times to state S_3 . Thus,

$$P(S_1|S_2, A_2) = \frac{300}{1000} = 0.3 \quad (3.7)$$

$$P(S_2|S_2, A_2) = \frac{200}{1000} = 0.2 \quad (3.8)$$

$$P(S_3|S_2, A_2) = \frac{500}{1000} = 0.5 \quad (3.9)$$

While both sets of transition parameters have the same value, the second set is more reliable. In order to take this lack of reliability into account, Tetreault and Litman proposed a CI estimate based upon the available data in [Tetreault and Litman, 2008], see also [Tetreault and Litman, 2006b, Tetreault and Litman, 2006a]. It is done by taking transition matrix T for slice and sample from each row using Dirichlet distribution for q times. As a result, it generates a large number of new transition metrics T_1, T_2, \dots, T_q that are all very similar to T . They then run MDP on all q transition matrices to get a range of ECR's (in this dissertation $m=1000$ was used, which is also used in [Tetreault and Litman, 2008]).

Their algorithm looks like this:

1. Compute transition probability matrix T from from the training data.
2. Use Policy iteration to compute an optimal policy π^* for S, A, T, R .
3. Sample q transition metrics T_1, T_2, \dots, T_q by sampling from the Dirichlet distributions corresponding to the counts observed in the training data;
4. Compute the value of the optimal policy π^* in each of these m models.
5. Numerically build the 95% confidence interval for the policy π^* based on the resulting value estimates: the bounds for the confidence interval are set at the lowest and highest 2.5 percentile of the resulting distributions.

3.6.3 An Example to Illustrate ECR and CI

This section illustrates ECR and CI with an example. In [Tetreault et al., 2007] the authors employed five feature choices to represent the state space: Certainty, Correctness, Percent Correct, Concept Repetition, and Frustration. Their system employed four possible tutor actions: ask a simple answer question; ask a complex answer question; ask a combination of the two; or do nothing. They estimated the reward value based upon the students' NLG. For the purpose of strategy induction, they assigned a reward of +100 if the students' NLG

π_1 : Features: Certainty, Correctness, and Concept Repetition. ECR: 42.56 95%CI: [28.37, 59.29] Interval Width: 23.52
π_2 : Features: Certainty, Correctness, and Percent Correctness ECR: 28.50 95%CI: [-5.89, 57.82] Interval Width: 63.71

Figure 3.1: ECR and CI Sample Learned policies

was above the median value. NLG scores below the median value were assigned a score of -100 . Two of the learned policies are summarized in Figure 3.1.

According to this assessment, π_1 will be both more effective and more reliable than π_2 because the former has a higher ECR, but a narrower CI than the latter. In Study 3, both the **ECR** and **95%CI** were employed as feature selection criteria. More specifically, the upper and lower bounds of the CI were used, and are referred to as the **lower-bound** and **upper-bound** of the policy.

3.7 ISSUE 5: CONFLICTING POLICIES

Once the tutorial tactics were induced, in order to test their effectiveness on the real subjects the researchers needed to implement them back to Cordillera. In order to execute these tutorial tactics, the dialogue manager needed to keep a record of the student's current states

on each KC. Moreover, it should also retain a KC-based record on the tutorial decision steps. So when a tutorial decision step occurred, the dialogue manager first looked up the KC(s) involved in that step and then looked up the corresponding policies. When a tutorial decision did not involve any specific KCs, the dialogue manager followed the KC-general tutorial tactics. When it involved a specific KC, the dialogue manager followed the tutorial tactics for that KC only. However, sometimes a tutorial decision involved multiple KCs, which generated conflicting decisions. In this case, the researchers needed to decide how the dialogue manager should deal with conflicting policies. This was the fifth and final issue requiring a decision.

3.8 DISCUSSION

To summarize, the general procedure for RL applications in this thesis was to first select a training corpus and then which KCs would be used to derive specific tutorial tactics for. For each KC_i , two types of tutorial decisions: $\langle ET, JS \rangle$ were derived and a KC-based reward R_{KC_i} for each student's tutorial dialogue is defined. The KC-based reward was defined based on the student's KC-specific NLG scores on KC_i . Then the KC-based feature choices were defined and each feature choice was discretized into discrete variables. Finally, the procedure described in Figure 3.2 was executed.

The five main RL-related issues addressed in this methodology are 1) selection of a training corpus, 2) choosing knowledge components to derive specific tutorial tactics on, 3) determining the state representation, 4) defining the reward function, and finally, 5) dealing with conflicting policies on multi-KC steps. State representation was divided into four sub-issues: a) defining feature choices, b) identifying feature discretization procedure, c) determining the feature selection procedure, and d) determining maximum number of features included in a policy.

In Chapters 5 and 6, the procedure for applying RL to derive KC-based tutorial tactics will be described, including how the five issues and four sub-issues were addressed. The rest of the procedure in both Studies 2 and 3 follows the general methodology described above.

1. Select representations for dialogue states S_{KC_i} , A_{KC_i} , and reward R_{KC_i}
2. FOR each subset of features S'_{KC_i} selected from S by following a feature selection procedure, do:
 - Use a training corpus to building an MDP' model based on S'_{KC_i} , A_{KC_i} and R_{KC_i} . The transition probability T'_{KC_i} of this MDP' is approximated based on the collected exploratory corpus.
 - Compute a policy from the learned MDP' by Tetreault, & Litman's toolkit.
3. Select a policy from all of induced policies that has the highest ECR.

Figure 3.2: General RL Procedure For Inducing KC-based Tutorial Tactics

This series of studies was designed to investigate two primary research questions: (1) Do pedagogical tutorial tactics on Elicit/Tell and Justify/Skip-justify impact students' learning? And (2) Is reinforcement learning a feasible method to induce tutorial tactics?

In Study 1, an exploratory corpus was collected by training a set of real students in a version of Cordillera that made random ET and JS decisions. In Study 2, we defined a set of 18 pedagogically relevant features, applied a greedy-like feature selection method to narrow the list down to four and applied RL to induce KC-based pedagogical tutorial tactics from the Exploratory corpus. The induced policies were then incorporated back into Cordillera and a second group of students was trained with this new version of Cordillera.

Finally, in Study 3, a set of 50 features was defined and a variety of feature selection methods were used to winnow them down to a set of six. More specifically, two sets of pedagogical tutorial tactics were induced: Normalized Gain (NormGain) and Inverse Normalized Gain (InvNormGain). The NormGain set was derived with the goal of enhancing the tutorial decisions that contribute to the students' learning; while the InvNormGain set was derived with the goal of enhancing those decisions that contribute less or even none to the students' learning. Both sets were then incorporated back into Cordillera, and students were trained

on the new versions of Cordillera with random assignments to conditions. The expectation was that the NormGain group would out-perform the InvNormGain group.

4.0 STUDY 1: EXPLORATORY CORPUS

The goal in Study 1 was to collect an exploratory corpus. The main advantage of collecting an exploratory training corpus is to potentially compute an effective policy within a large state space using a relatively small amount of training data. It addresses the situation in which collecting real-world experience is highly costly, but computation is relatively cheap. For example, it has been shown that this approach is effective at automatically learning the effective action to take in any state in various dialogue systems where collecting data is even less expensive than ITSs [Williams et al., 2005, Walker, 2000, Singh et al., 2002].

Study 1 used the Random-Cordillera on which the dialogue manager made random decisions at each tutorial decision step. A set of real human participants interacted with Random-Cordillera, from which an exploratory training corpus was collected for deriving pedagogical tutorial tactics in Study 2 and Study 3. In the sections below this process will be described in detail.

4.1 METHODS

4.1.1 Participants

Data was collected over four months during Fall 2007. Seventy college students were recruited. They were required to have a basic knowledge of high-school algebra, but not to have taken college-level physics courses. All subjects were paid for their time, regardless of completion. Subjects who completed the study took from two to three weeks to complete the study. In all, 64 students completed the experiment.

4.1.2 Materials

As with the other studies in this thesis, Study 1 was done in the Physics work-energy domain, a common component of introductory college physics courses.

4.1.2.1 32 Knowledge Components Two domain experts (not the author) who are also knowledge representation experts, identified 32 KCs in the domain (see Appendix A). They had experience identifying KCs for a series of previous studies involving college physics. One example of their work is KC_{20} which is defined as:

Definition of Kinetic Energy ($KE = \frac{1}{2}mv^2$) — KC_{20} : If an object is moving, then its kinetic energy at a time is $\frac{1}{2}mv^2$, where m is the object’s mass and v is the magnitude of the object’s instantaneous velocity.

4.1.2.2 Physics Textbook The physics textbook used in this study is web-based. It was written by a domain expert who is also a native English speaker (not the author). It includes all the physics concepts that were needed for the domain. For each physics concept and domain principle, a general description was presented together with some worked examples (see Appendix D). For example, the description of KC_{20} in the textbook begins with “One type of energy, called kinetic energy (KE), is associated with individual objects. It depends only on an object’s mass and on the magnitude of its velocity...”. More information can be found at section D.0.2.13 in Appendix D. The textbook was 27 pages long. When reading the textbook, students were free to move forward and back. Appendix D shows the textbook content.

4.1.2.3 Pre- and Posttest The pre- and post-tests were identical in Study 1. Both contained a total of 33 problems selected from the Physics literature (see Appendix E) by two domain experts (not the author). The tests were given online and consisted of both multiple-choice and open-ended questions. The latter questions required the students to derive an answer by writing or solving one or more equations. Once an answer was submitted, the students automatically proceeded to the next question without receiving any feedback on the

Table 4.1: Major Principles of Work and Energy

KC	Descriptions of the principles	Expressions
KC_1	Weight Law (w)	$W = mg$
KC_{14}	Definition of Work (W)	$W = Fd\cos(\alpha)$
KC_{20}	Definition of Kinetic Energy (KE)	$KE = \frac{1}{2}mv^2$
KC_{21}	Gravitational Potential Energy (GPE)	$GPE = mgh$
KC_{22}	Spring Potential Energy (SPE)	$SPE = \frac{1}{2}kd^2$
KC_{24}	Total Mechanical Energy (TME)	$TME = KE + GPE + SPE$
KC_{27}	Conservation of Total Mechanical Energy (CTME)	$TME_1 = TME_2$
KC_{28}	Change of Total Mechanical Energy for Non-isolated Systems (TMENC)	$Net_W = TME_2 - TME_1$

correctness of a response. Students were not allowed to return to prior questions. Appendix F listed the number of times each KC showed up in the tests. Except for KC_7 , all the rest of the KCs appeared in at least one test item. For example: the first test problem is an open-ended question involving KC_{20} . It stated:

1. Enter the equation that defines the kinetic energy of an object (remember to use * for multiplication and ^ for exponentiation, and be sure to include an = sign):

4.1.2.4 Domain Principles The eight major principles in the domain are shown in Table 4.1. In Table 4.1, the first column lists its corresponding KC number. The second column describes the name of the principle. The last column is the formula or mathematical expression of the principle. To differentiate these KCs from the rest of 24 non-domain principle KCs, the name the domain principle-related KCs were named as *primary KCs*. As the table shows there are eight primary KCs in this domain. For example, the fourth row in Table 4.1 is the definition of Kinetic Energy and its corresponding KC is KC_{20} .

4.1.2.5 Seven Training Problems Participants solved a series of seven training problems. The problem statements are listed in Appendix G. The problems were arranged in order of increasing complexity. Table 4.2 contains a list of the problems in the order they were presented to the students and identifies which of the eight main KCs were relevant to each problem. For example, P4 is an example used earlier in this dissertation. It is defined as follows:

A 0.6kg rock in space has a velocity of magnitude 2.0m/s at point A and kinetic energy of 7.50J at point B. What is the net work done on the rock as it moves from A to B? We define:

T_0 : the time point when the rock is at point A.

T_1 : the time point when the rock is at point B.

As mentioned in an earlier chapter, solving training problem P4 involved applying three major domain principles, with some principles needing to be applied twice. The three domain principles are: the definition of Kinetic Energy (KE: $KE = \frac{1}{2}mv^2$), the definition of Total Mechanical Energy (TME: $TME = KE + GPE + SPE$), and the Change of Total Mechanical Energy for Non-isolated Systems ($NetW = TME_2 - TME_1$). These were represented as KC_{20} , KC_{24} , and KC_{28} respectively. Therefore, the fifth row in Table 4.2 shows that the relevant KCs for the training problem P4 are KC_{20} , KC_{24} , and KC_{28} .

Table 4.2: Seven Training Problems

Problems:	Primary KCs							
	KC_1	KC_{14}	KC_{20}	KC_{21}	KC_{22}	KC_{24}	KC_{27}	KC_{28}
P1			X					
P2		X						
P3					X			
P4			X			X		X
P5			X	X		X	X	
P6	X		X	X		X	X	
P7			X	X	X	X	X	

4.1.2.6 Training Scripts For each training problem, a tutorial script was written to cover the relevant content. The content included how to solve the problem and the post-problem discussions. The tutorial scripts were written collaboratively by a group of five

authors (including this author). Four of the five were domain experts, while one was trained in tutoring, specifically in the authoring of tutoring dialogues. Two of the domain experts were native English speakers while the remainder were fluent in English. Appendix I provides an example of the written tutorial script for training problem P4. Eventually each script was checked by at least two out of five domain experts, who agreed upon the content of the script.

4.1.2.7 Random-Cordillera The students in Study 1 trained on Random-Cordillera. The dialogue manager on Random-Cordillera did not follow any tutorial tactics, but made all the tutorial decisions randomly.

4.1.2.8 Human Wizards As described earlier, in order to reduce confounds due to imperfect Natural Language understanding, human wizards were used. In Study 1, there were a total of six human wizards. Their sole function was to map students' entries to the closest answer. They cannot control which tutorial actions the dialogue manager should take next.

4.1.2.9 Some Clarification On The Number Of KCs Appearing In This Dissertation As mentioned before, for the Work and Energy domain, we have identified a total of 32 KCs. Among them, 31 KCs were evaluated in the pre- and post-tests (KC_7 was not). Of these 31 KCs, 21 KCs were involved in the ET decisions while 10 KCs were involved in the JS decisions. In Study 2, KC-specific tutorial tactics for all possible KCs were learned, so 21 KC-specific ET tutorial tactics and 10 JS ones were induced. In Study 3, however, the main focus was on the eight *primary* KCs that are also domain principles. Among the eight primary KCs, KC_1 did not show up in any JS tutorial decisions. Therefore, in Study 3, eight KC-specific ET tutorial tactics and seven KC-specific JS tutorial tactics were induced.

4.1.3 Procedure

The study itself consisted of five standard phases: 1) background survey, 2) pre-training, 3) pre-test, 4) training, and 5) post-test. In each phase, students were not restricted to any

time limits. This was also true for Studies 2 and 3. The background survey asked students for demographic information such as gender, age, SAT scores, high school GPA, experience with algebra, calculus, physics, and other information (see Appendix C). Following the background survey, students read the physics textbook during the pre-training (see Appendix D) and took the pre-test (see Appendix F). The physics textbook was only available during phase 2, pre-training. This was also true for Studies 2 and 3.

In phase 4, students were first trained to solve a demonstration problem, which did not include physics content, on Cordillera. The sole purpose of this step was to familiarize them with the GUI interface. They then solved the seven training problems on Random-Cordillera. Finally, students took the post-test which was identical to the pre-test. Students were given the same set of questions in the same order.

4.1.4 Grading

All of the tests were graded in a double-blind manner by a single domain expert who was an experienced grader (not the author). In a double-blind manner, neither the students nor the grader know who belongs to which group. In this case, the grader was not familiar with the hypotheses being tested. Each test question was assigned two grades: overall and KC-based grade. The overall grade was a score in the range $[0, 1]$ describing the correctness of an answer as a whole. Since there were 33 test questions in all tests across the three studies, the maximum overall score for each test was 33.

Under the KC-based grading criteria, the grader first identified all of the relevant KCs for a test question, and then assigned a score in the range $[0, 1]$ for each KC application. Each of these scores was called the KC-based score. Except for KC_7 , all the remaining 31 KCs were present in at least one question of the tests. The maximum score for a test question under the KC-based grading was the number of KCs involved in the question. A KC cumulative score was calculated for each student by summing up all KC-based scores across all of the test questions in the test. In this thesis, there were a total of 168 KCs in all pre- and post-tests.

For each KC, e.g. KC_k , a KC-based score was calculated by simply summing over the

KC-based scores on KC_k across all 33 test questions divided by the number of test questions problems that involves KC_k in the test. The number of occurrences of each KC in the pre- and post-tests in this dissertation is shown in the last row in Appendix F. It shows that the frequencies of 31 KCs in the test vary from one to up to 16.

The following example illustrates these three grading scores. Figure 4.1 presents a student's answer to a test question 10. In this example, the question statement is listed first and then the correct answer. The mapped KCs refer to the two KCs that needed to be applied to solve the question: KC_{20} and KC_{23} . Their corresponding descriptions are also listed. The student's answer is partially correct in that the number is correct but without unit. So the overall score for this answer is 0.75 (maximum is 1 for each test question) and the cumulative KC score is 1 (the maximum is 2 which is the number of KCs involved in the test question 10). Additionally, the student also received two KC-based scores for this problem: 1 for KC_{20} and 0 for KC_{23} .

For comparison purposes all of the scores were normalized to fall in the range of [0,1]. Most of the analysis in the following sections is based upon the overall and cumulative KC scores. The KC-based scores will be presented only for Study 1 and Study 2. This is because the KC-based scores are used to describe the characteristics of the Exploratory and the Dichotic Gain (DichGain) Corpora and the KC-based NLGs (calculated from the KC-based pre- and post-test scores) were used to derive KC-based tutorial tactics in Studies 2 and 3.

4.1.5 Measures

The main purpose of Study 1 was to collect an Exploratory corpus. Therefore, the focus will primarily be on two aspects: learning performance and the characteristics of the Exploratory corpus. For the learning performance, students' pre- and post-test scores were compared under both grading criteria to determine whether the Exploratory group learned by training on Random-Cordillera.

For the characteristics of the corpus, the average number of ET and JS decisions and overall decisions across students' tutorial logs were presented. The I-ratio and the J-ratio

Question 10: A toy cart moves with a kinetic energy of 30 J. If the magnitude of its velocity is doubled, what will its kinetic energy be?

Correct Answer:

1. Kinetic energy = 30J
2. $KE = 1/2 * m * v^2$
3. $new_v = 2 * v$
4. $new_{KE} = 1/2 * m * (2v)^2 = 4 * (1/2 * m * v^2) = 4 * KE = 4 * 30J = 120J$

Mapped KCs:

KC_{20} : definition of kinetic energy $KE = 1/2 * m * v^2$

KC_{23} : The unit for energy is the Joule (J)

Student Answer: $new_{KE} = 1/2 * m * (2v)^2 = 4 * (1/2 * m * v^2) = 120$

Overall Score : 0.75 (maximum is 1 for this problem)

Mapped KC Score: $\{KC_{20}: 1\}; \{KC_{23}: 0\}$

Culmulative KC-based Score : 1 (maximum is 2 for this problem)

Figure 4.1: An Example of Three Grading Criteria

were also checked to determine whether the random decisions worked. Recall that previously the I-ratio and J-ratio were defined as:

$$\mathbf{I - ratio} = \frac{N_{Elicit}}{N_{Elicit} + N_{Tell}} \quad (4.1)$$

$$\mathbf{J - ratio} = \frac{N_{Justify}}{N_{Justify} + N_{SkipJustify}} \quad (4.2)$$

The higher the I-ratio is, the more interactive the dialogue might be. The higher the J-ratio is, the more likely the students would be presented a justification step. Specifically, the average number of justification steps each student received was also presented. Additionally, because the Exploratory Corpus would be used to induce KC-specific tutorial tactics in

Studies 2 and 3, the analysis will include discussion of the characteristics of the Exploratory Corpus and the KC-based learning performance.

Study 1 was conducted as a part of an NSF ITR (Information Technology Research)¹ project, for which this author acted as a script author developing tutorial scripts for the training problems. This author was primarily responsible for the scripts associated with training problems P4 and P5. The scripts were checked by other members of the group. This author was also involved in annotating training problems with relevant KCs and acting as a human wizard during the collection of the Exploratory Corpus.

4.2 RESULTS

4.2.1 Time

Each student took between six and fourteen hours (3-7 sessions) to finish the study. Each session typically lasted about two hours. In general the students spent roughly five hours (ranging from as low as four to as high as nine hours) training with Random-Cordillera. For analysis purposes, each student's training dialogues were concatenated into a single consecutive dialogue resulting in a single "super-dialogue" for each student. These super-dialogues, together with pre- and post-test KC-based scores were used to derive the KC-based policies for use in Studies 2 and 3.

On average, it took each student $M = 280.38$ mins ($SD = 66.88$) to finish the seven training problems. For each training problem, the average time spent by the Exploratory Group varied. In Figure 4.2, the y-axis shows the time on task in minutes. It shows the Exploratory students spent less than 30 minutes on each of the first three simple problems, while on the last problem P7, they averaged 67 minutes, the longest average time per problem. As the training problems became more complicated, the more time it took the Exploratory students to finish it.

¹Support for this research was provided by NSF grants #0325054

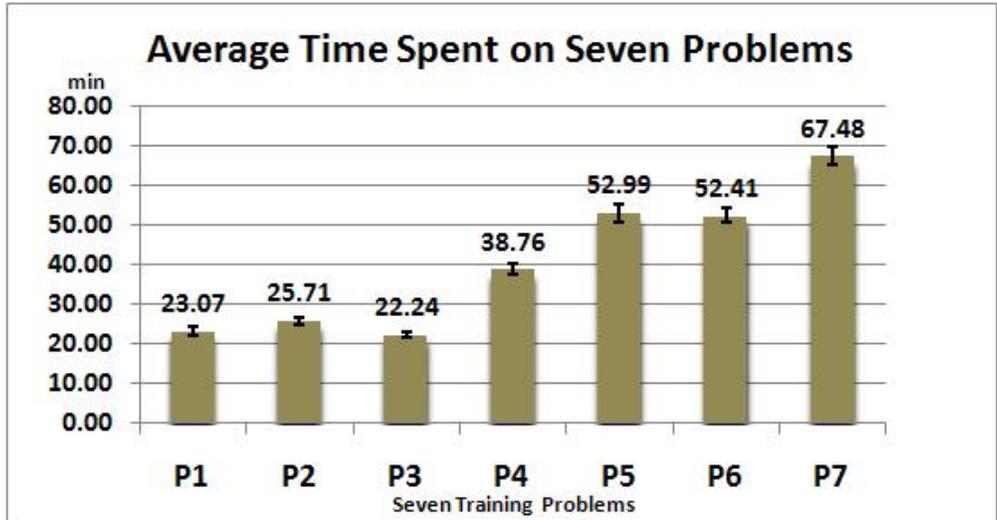


Figure 4.2: Average Time Spent Per Training Problem By the Exploratory Group

4.2.2 Learning Results

A one-way ANOVA was used to test for performance preference differences between the pre- and posttests. Regardless of grading criteria, participants made significant gains from pre-test to post-test (Table 4.3), $F(1, 126) = 18.76, p < .001, R^2 = .73$ under the overall grading criteria and $F(1, 126) = 11.01, p = .001, R^2 = .69$ under the cumulative KC-based grading criteria. Table 4.3 also lists the overall NLG scores and the Cumulative KC-based NLG scores. Since a student's NLG is defined as his or her actual average learning gain divided by his or her maximum possible actual gain, on average the Exploratory students made 29% gains under the overall grading criteria and 25% of the possible gains under the cumulative KC scoring rubric.

Figure 4.3 shows the Exploratory group's pre- and post-test scores under the two grading criteria. A double asterisk (**) shows that the difference is statistically significant ($p < 0.05$). Table 4.3 summarizes the minimum, maximum, mean, and SD of each score.

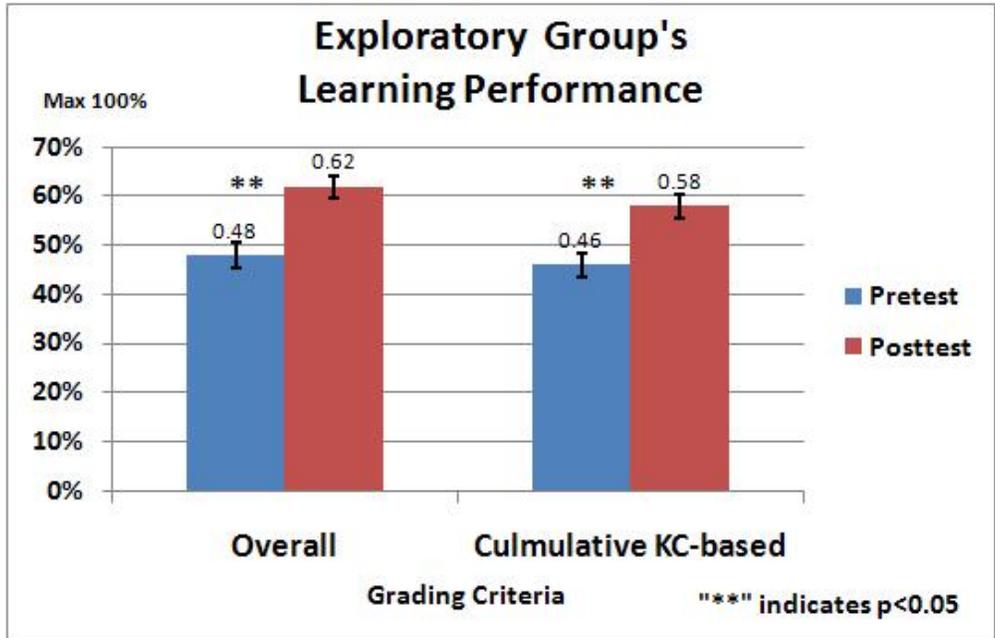


Figure 4.3: Learning Performance on Exploratory Group

Table 4.3: Exploratory Students' Learning Performance

Grading		Min	Max	Mean	σ
Overall	Pre-test	0.14	0.94	0.48	0.20
	Post-test	0.21	1.00	0.62	0.18
	NLG	-0.08	1.00	0.29	0.21
Cumulative KC-based	Pre-test	0.13	0.96	0.46	0.20
	Post-test	0.11	1.00	0.58	0.20
	NLG	-0.38	1.00	0.25	0.26

4.2.3 Exploratory Corpus

4.2.3.1 Overall Characteristics The total number of ET tutorial decisions, referred to as *ET decisions* in this thesis, ranged from 250 to 332 ($M = 273.89, SD = 12.46$) and that of JS tutorial decisions, referred to as *JS decisions*, ranged from 52 to 71 ($M = 56.61, SD = 3.43$). The total number of the tutorial decisions regardless of actions, referring to as *overall decisions*, for each student ranged from 288 to 372 ($M = 305.48, SD = 14.01$)².

In the Exploratory Corpus, since all tutorial decisions were randomly decided, the dialogues' I-ratio was expected to be around 0.5. An analysis of the log files showed that the I-ratios ranged from 0.44 to 0.56 ($M = 0.50, SD = 0.03$). Similarly, it was expected that the J-ratio would be roughly 0.5 as well. The analysis of the log files showed that this value ranged from 0.39 to 0.68 ($M = 0.53, SD = 0.06$). The larger justification range is unsurprising as there were fewer possible justification steps in the script. More specifically, the number of justification steps in a student's tutorial dialogue ranged from 21 to 40 ($M = 30.17, SD = 3.83$).

Table 4.4 summarizes the overall characteristics of the tutorial decisions in the Exploratory Corpus.

Table 4.4: Overall Characteristics On Tutorial Decisions in Exploratory Corpus

Value	Min	Max	Mean	σ
ET Decisions	250	332	273.89	12.46
JS Decisions	52	71	56.61	3.43
Overall Decisions	288	372	305.48	14.01
I-ratio	0.44	0.56	0.50	0.03
J-ratio	0.39	0.68	0.53	0.06
Justify	21	40	30.17	3.83

²overall decisions < ET decisions+JS decisions because on certain tutorial decision steps, the tutor makes both types of decisions: JS first and then ET. For instance: Line 3 in Figure 2.1. When we calculated the overall decisions, such a step was counted as one decision step.

4.2.3.2 KC-based Characteristics As described above, not all KCs appeared on tutorial decision steps in the authored scripts. In order for the tactic induction process to be effective, a KC must be involved in the tutorial dialogues, the decision steps and the pre- and post-tests. There were a total of 21 KCs that satisfied this requirement. All 21 KCs appeared in at least one ET decision step. Only 10 out of 21 KCs appeared in at least one JS decision step. The characteristics of the 21 KCs in the Exploratory Corpus will be discussed next.

Table 4.5 presents KC-based scoring and breakdown for the KC-related information in the corpus. More specifically, the second column lists the 21 KCs. The third and fourth columns list the average number of ET and JS decisions for the corresponding KC respectively. The last column presents the comparison of KC-based post-test and pre-test scores to show whether students' performance on the KC was improved after the seven training problems.

From columns 2, 3 and 4, we can see that the number of decision steps varies dramatically across the KCs. Column 2 shows that the average number of ET decisions ranges from 1.5 for KC_{18} to 72.6 for KC_{20} : the definition of kinetic energy ($KE = \frac{1}{2}mv^2$); column 3 shows that 10 KCs appeared in the JS decision steps and the average number of JS decisions ranges from 2 for KC_{12} , KC_{14} , KC_{26} to 16.9 for KC_{21} : definition of Gravitation Potential Energy ($GPE = mgh$). The average number of overall decisions varies from 2 on KC_{18} to 81.5 for KC_{20} .

In Table 4.5, the last column shows that students learned significantly on 13 out of 21 KCs (labeled with “**”) and on the remaining KCs, no significant difference was found between their pre- and post-test scores.

The eight primary KCs are underlined. The Exploratory Group scored significantly higher on the posttest than on the pre-test for six of the eight KCs. The two exceptions are KC_{14} and KC_{28} . Among the eight primary KCs, KC_1 never occurred in any of the JS decisions. The remaining seven KCs appeared in both ET and JS decisions.

Table 4.5: KC-based Exploratory Corpus

	KC	ET	JS	Overall Decisions	Compare Pre- and Posttests
1	<u>KC₁</u> **	4		4	$t(126) = 3.280, p = 0.001$
2	KC ₃	2.1		2.1	$t(126) = 1.204, p = 0.231$
3	KC ₅	6.9		6.9	$t(126) = 1.257, p = 0.211$
4	KC ₉ **	2		2.5	$t(126) = 3.527, p = 0.001$
5	KC ₁₂	2.1	2	3	$t(126) = 1.392, p = 0.166$
6	KC ₁₃	3		3	$t(126) = 1.560, p = 0.121$
7	<u>KC₁₄</u>	8	2	9	$t(126) = 1.076, p = 0.284$
8	KC ₁₅ **	7.5		7.5	$t(126) = 2.470, p = 0.015$
9	KC ₁₇	3.8		4.4	$t(126) = .880, p = 0.381$
10	KC ₁₈	1.5		2	$t(126) = 1.478, p = 0.142$
11	<u>KC₂₀</u> **	72.6	15.4	81.5	$t(126) = 5.379, p = 0.000$
12	<u>KC₂₁</u> **	33.6	16.9	60.8	$t(126) = 3.932, p = 0.000$
13	<u>KC₂₂</u> **	30.7	3.3	32.1	$t(126) = 2.389, p = 0.018$
14	KC ₂₃ **	62.1	4	63	$t(126) = 5.358, p = 0.000$
15	<u>KC₂₄</u> **	52.7	15.5	60.5	$t(126) = 3.924, p = 0.000$
16	KC ₂₅ **	8.7		9.3	$t(126) = 3.767, p = 0.000$
17	KC ₂₆ **	4.2	2	6.1	$t(126) = 2.063, p = 0.041$
18	<u>KC₂₇</u> **	21.5	4.8	23.9	$t(126) = 4.522, p = 0.000$
19	<u>KC₂₈</u>	14.2	4.1	16.6	$t(126) = 1.911, p = 0.058$
20	KC ₃₁ **	18.6		19.2	$t(126) = 2.446, p = 0.016$
21	KC ₃₂ **	14.4		14.4	$t(126) = 3.888, p = 0.000$

4.3 DISCUSSION

The goal of Study 1 was to collect an Exploratory Corpus. Ideally, the Exploratory Corpus should be collected by exploring all possible states and testing all possible actions from each possible state, preferably for many times. But given the high cost of collecting these educational data, it is not possible to do so. Therefore, in this thesis the Exploratory Corpus was collected by choosing actions randomly. It was expected that viable, effective tutorial tactics would be collected from the “Exploratory” training corpus. Our analysis of the corpus showed that random decisions seemingly balanced the number of elicits and tells students got during the tutoring (The mean of the I-ratio was 0.50). It was less so for JS decisions, but the mean was off only slightly. (The mean of the J-ratio is 0.53). It also demonstrated that the ET decisions were four times more frequent than the JS decisions. Moreover, 21 KC-based ET tutorial tactics and 10 KC-based JS ones could be induced from the Exploratory Corpus.

While randomness may not be the best guide, the students made significant learning gains in Study 1 as evidenced by their pre- and post-test scores. The results seemed to confirm the previous view that content exposure might cause students to learn even from tutors with non-optimal pedagogical skills.

5.0 STUDY 2: DICHOTIC GAIN (DICHGAIN) GROUP

The main goal of Study 2 was to investigate whether the induced tutorial tactics from the Exploratory Corpus would result in a more effective version of Cordillera. Ideally, in order to investigate the effectiveness of the system, a full-scale comparison between the new system and Random-Cordillera should be conducted by randomly assigning students to one of two groups. However, given the cost of running a comparison, and the issues to be addressed in order to apply the RL to induce tutorial tactics, Study 2 was treated as an engineering project rather than a science project. A new group of subjects were tested on the new system and the students' results were compared with the Exploratory group. This type of comparison is not rare and is used if the goal is to determine if a trend exists. For example, in [Singh et al., 2002] the researchers tested the learned policy on a new group of users and compared the new group's results with the original training group.

If there was a trend showing the new system out-performing the initial Random-Cordillera, a full-scale comparison would be conducted. However, a subsequent analysis suggested that the learned policies may not be very effective. This led to the hypothesis that this trend might be caused by the limited methodology used for applying RL to induce tutorial tactics. For example, a greedy-like feature selection method was used to derive a set of pedagogical tutorial tactics from the Exploratory Corpus collected in Study 1 (described below). Nevertheless, an important contribution of Study 2 is that it generated a new corpus that is similar to many other preexisting corpora and datasets in that it follows some types of tutorial tactics, whether effective or not. This new corpus can be further used to induce new tutorial tactics. If a successful policy can be induced from this corpus, then it will show the potential for applying RL repeatedly to improve ITSs from pre-existing data.

Based on the reward functions employed in Study 2, the induced tutorial tactics were referred to as Dichotic Gain (DichGain) tutorial tactics and the new version of the system was labelled DichGain-Cordillera. The new group of students who were trained on DichGain-Cordillera was named the DichGain group, and the new corpus was labelled the DichGain Corpus. This chapter will describe Study 2 in detail.

5.1 APPLY RL TO INDUCE DICHGAIN POLICIES

Recall that the tutorial actions which are the central focus of this thesis were elicit/tell (ET) and justify/skip-justify (JS). Chapter 3 provided a general overview of Reinforcement Learning (RL) and described the toolkit used to derive tutorial tactics in this dissertation. It was assumed that KC-based tutorial tactics would be more effective than KC-general ones. Five issues needed to be addressed to induce KC-based tutorial tactics for each type of action. The issues described were: 1) obtaining a training corpus; 2) selecting the target KCs for which the tactics will be induced; 3) defining the reward function, 4) determining state representation; and finally, 5) selecting a conflict-resolution policy for multi-KC decision steps. In this section, it will be shown how those issues were addressed in Study 2.

5.1.1 Training Corpus

This study made use of the Exploratory Corpus collected in Study 1. That corpus consisted of sixty-four student tutorial dialogues, one for each participant. Each dialogue covered the entire interaction between the student and the Random-Cordillera system over the seven training problems.

5.1.2 Knowledge Components

In order for the tactic induction process to be effective, a KC must be involved in the tutorial dialogues, the decision steps and the pre- and post-tests. Table 4.5 in Chapter 4 lists the 21 KCs that meet these criteria. For this study, the decision was to induce one policy for

each relevant KC. Therefore, of the thirty-two KCs in the domain, a total of thirty-one KC-based tutorial tactics were induced: 21 KC-based tutorial tactics on ET decisions and ten KC-based tactics on JS decisions. Additionally, two KC-general tactics, one for ET decisions and another for JS decisions were induced. The KC-general tutorial tactics were used in the decision steps that did not involve any of the identified KCs. In short, this resulted in a total of thirty-three tutorial tactics: twenty-two ET tactics and eleven JS tactics. In the following discussion, the study uses $\pi(KC_i, ET)$ and $\pi(KC_i, JS)$ to refer to KC-based tutorial tactics for KC_i for the ET and JS decisions respectively. For KC-general tutorial tactics, the study uses $\pi(KC*, ET)$ and $\pi(KC*, JS)$.

5.1.3 KC-based Reward

As described in Chapter 3, for a tutorial dialogue d_i on KC_k , there are a set of intermediate rewards for each state and KC: $r_{d_i, kc_k}^1 \dots, r_{d_i, kc_k}^{n_{d_i, kc_k}-1}$ all of which are equal to 0. Only the final reward, $r_{d_i, kc_k}^{n_{d_i, kc_k}}$, has a non-zero value. Here n_{d_i, kc_k} represents the number of dialogue turns d_i in which the system made decisions regarding KC_k .

In Study 2, a similar approach to reward function was used as [Tetreault and Litman, 2008]. The student’s final reward for each KC_k in his/her superdialogue d_i was based upon his/her KC-based NLG for KC_k . More specifically, for each KC_k , the students were divided into two groups, *low learners* and *high learners*, according to a median split of the students’ KC-based NLGs.

The high learners were assigned a final reward of +100, while the low learners were assigned a final reward of -100. These final reward values will be propagated to the internal states via a dynamic programming algorithm for policy iteration [Sutton and Barto, 1998]. For inducing KC-general tutorial tactics, the reward functions were based on the cumulative KC-based NLG instead of KC-based NLG on a specific KC. The rest of the procedure remained the same.

5.1.4 State Representation

As described in Chapter 3, the issue of the state representation can be divided into four sub-issues. They are: 1) defining the potential feature choices in state representation; 2) capping the number of features included in each policy; 3) discretizing the features appropriately; and 4) determining feature selection procedures. The next section describes how these sub-issues were addressed in Study 2.

5.1.4.1 Feature Choices

Moore et al. identified four types of features that are relevant for human tutors when making tutorial decisions: autonomy, temporal situation, problem-solving state, and performance to model the dialogue and the student’s state [Moore et al., 2004]. For each category three to seven features were defined. Note that in this dissertation only features that could be both automatically computed and unambiguously evaluated were included. This was because the tutor would require the features to be available in real time when the learned policies were employed. In order to help readers to understand each feature better, at the end of this subsection I will use a simplified example to illustrate how these features were calculated in Study 2.

Autonomy — three features

Autonomy features are related to the amount of work performed by the student in the dialog. All autonomy features end with an ‘A’ in their name and are numeric. They are

- F1 tellsSinceElicitA:** The number of tells the student has received since the last elicit prompt, irrespective of the KC involved. For example, $\text{tellsSinceElicitA} = 2$ means that two tell decisions have been made since the last elicit decision. This feature reflects how active a student is currently, that is, how much work the student has performed recently.
- F2 pctElicitA:** The percentage of elicit/tell decision points in which the tutor has opted to elicit during the dialog, irrespective of KC. This feature describes how interactive the overall tutorial dialogue is. If answering questions makes the student more active and interactive than simply receiving information from the tutor, then the higher the value of pctElicitA is the more active and interactive the tutorial dialogue is.
- F3 pctTellsKCSessionA:** The percentage of tells received in a session for a specific KC, e.g. KC_{20} . This feature describes how interactive the tutorial dialogue is for this session. This feature measures the autonomy characteristics of the student’s tutorial dialogue. It uses a longer timeframe than tellsSinceElicitA , but a smaller one than pctElicitA .

Of these features, `pctTellsKCSessionA` (F3) is KC-specific. The focus is on the value specific to the current KC. For example, if a policy for KC_{21} is induced, this feature computes the KC performance in terms of the tutorial actions on all previous instances of KC_{21} solely. In order to differentiate from other feature choices, a label “KC” was added to the name for all KC-specific features.

Temporal Situation — 3 features

Temporal situation features cover temporal information such as time spent on the training thus far. All three temporal situation features end with a ‘T’ and were numeric. Three features were defined:

- F4 `durationKCT`:** Time duration since the last tutorial decision was made on the current KC. e.g KC_{20} . The feature reflects how active a student’s knowledge of the current KC is. For example, if “`durationKCT`” is high, it means that the tutor has not mentioned the KC recently, so the student’s knowledge on the current KC may be still.
- F5 `TimeInSessionT`:** The total time in the session so far. This feature can be used to measure the student’s fatigue level.
- F6 `TimeBetweenSessionsT`:** The time elapsed between the end of the previous session and the beginning of the current one. This feature reflects how likely it is that a student has forgotten what was learned in previous sessions. The higher the value, the more likely the student has forgotten what was previously learned. If `TimeBetweenSessionsT` is high, then the tutor should probably remind the student of some domain knowledge at the beginning of the session.

Among them, `durationKCT` is a KC-specific feature.

Problem Solving Context — 5 features

Problem solving features include state information, such as what phase the dialogue is in (e.g. problem solving or post-problem discussion), the problem’s difficulty level, and so on. All problem solving-related features end with “PS.” Five feature choices are defined below.

- F7 `EarlyTrainingPS`:** Problems P1, P2, P3 are categorized as early training problem and the rest four training problems are categorized as late ones. For early training problems, the tutor may ask students to practice certain entries to let them get used to the tutor.
- F8 `ProblemComplexityPS`:** Problems P1, P2, P3 are simple; P4 and P5 are medium; and P6 and P7 are complex. This feature reflects the increasing complexity of the solutions for training problems. The feature is relevant because it is expected that

fewer students will be able to solve a training problem on his/her own as the complexity of the training problem increases.

- F9 DuringWalkThroughPS:** For each training problem the tutorial dialogues follow a two-phase procedure: first problem solving, followed by a post-problem discussion. This feature describes whether a tutorial decision was made during the problem solving or during the post-problem discussion. It is probably relevant to learning because certain tutor actions would be better to happen during the problem solving but not post-problem discussion and vice versa. For example, qualitative discussions may sometimes distract students from problem solving, and thus it would be better to reserve qualitative discussion for the post-problem discussion rather than during the problem solving itself.
- F10 nKCsPS:** The number of times the current KC has occurred in the current tutorial dialogue. This feature reflects the student’s overall familiarity with the current KC.
- F11 nKCsSessionPS:** The number of occurrences of the KC, e.g. KC 20 in this session. This feature reflects how many times the student has accessed the current KC in this session.

Two features are KC-specific: nKCsPS(F10) and nKCsSessionPS (F11).

Performance — seven features

Performance features describe factors such as the quality of the student’s previous answers and the student’s ability. All performance-related features end with “PM”. Seven have been defined. These are described below:

- F12 pctCorrectPM:** Defined as: $\frac{\text{correct}}{\text{correct}+\text{incorrect}}$ on all KCs. The number of *correct* and *incorrect* entries calculated in the students’ logs that were labelled with + and – respectively. This feature measures the student’s overall competence when only elicits are counted as learning opportunities.
- F13 pctOverallCorrectPM:** Defined as $\frac{\text{correct}}{\text{correct}+\text{incorrect}+\text{tells}}$ on all KCs. This feature is probably relevant to learning in that it reflects the student’s overall competence when both elicits and tells are counted as learning opportunities.
- F14 nCorrectKCPM:** The number of correct responses on the current KC, e.g. KC 20. This feature reflects the student’s competence on the current KC.
- F15 pctCorrectKCPM:** Defined as $\frac{\text{correct}}{\text{correct}+\text{incorrect}}$ on the current KC. e.g. KC_{20} . This feature is probably relevant to learning in that it reflects the student’s competence on the current KC when only elicits are counted as learning opportunities.
- F16 pctOverallCorrectKCPM:** Defined as $\frac{\text{correct}}{\text{correct}+\text{incorrect}+\text{tells}}$ on the current KC. e.g. KC_{20} . This feature reflects the student’s competence on the current KC when both elicits and tells are counted as learning opportunities.
- F17 nIncorrectKCPM:** The number of incorrect responses on the current KC, e.g. KC 20. This feature reflects the student’s incompetence on the current KC.
- F18 pctCorrectKCSessionPM:** Defined as $\frac{\text{correct}}{\text{correct}+\text{incorrect}}$ on the current KC. e.g. KC_{20} in this session. This feature reflects the student’s lack of competence on the current KC in this session.

Five out of seven features are KC-specific. They are: `nCorrectKCPM(F14)`, `pctOverallCorrectKCPM(F15)`, `pctCorrectKCPM(F16)`, `nIncorrectKCPM(F17)`, `pctCorrectKCSessionPM(F18)`.

As mentioned above, a successful application of RL is heavily dependent upon choosing an appropriate set of features to represent tutorial contextual states. In other words, the state representation for RL should include all of the tutorial dialogue information that is relevant and necessary to determine what action should be taken next. On the other hand, user modeling focuses on developing cognitive models of human users, such as the modeling of users' skills, knowledge level, and so on. Therefore, the features included in the state representation for the RL should include, but not be limited to, the features that model human users. In this project, some features, especially performance-related features as defined in Study 2 and Study 3 (described in Chapter 6) can be seen as modeling students' knowledge levels. One example of such a feature is “`pctCorrectKCPM`” which represents the percentage of times a student had the correct answer on a specific KC.

Earlier an explanation was provided for how KC-specific features were calculated for inducing specific KC tutorial tactics. However, when inducing KC-general tutorial tactics on either ET or JS decisions, all the KC-specific features become KC-general features and take into count all of the previous instances regardless of KC. For example, `nCorrectKCPM` becomes the number of correct responses on all the KCs instead of on a specific KC.

This section explains how the eighteen features were calculated from the log files. The following is an sample of a tutorial dialogue that was extracted from log files of a student solving the training problem P4. The entire dialogue between the student and RandomCordillera is contained in Appendix H. All the tutor turns and the student turns are labelled to the corresponding KCs in Appendix H. The sample dialogue shown here covers one step applying KC_{20} : the definition of Kinetic Energy, to solve the KE of the rock at T_0 .

The sample dialogue covers five micro-steps. They represent the first principle application in solving the training problem P4. So $EarlyTrainingPS = 0$ and $ProblemComplexityPS = medium$. The dialogue occurred during problem solving as opposed to post-problem discussion so $DuringWalkThroughPS = 1$. The sample dialogue happened in the student's first

training session on Cordillera, so the time duration between his current session (start time of the current session) and his last session (the end time of last session) $TimeBetweenSessionsT$ is 0. In fact, this is the student’s fourth training problem in the current session.

In Table 5.1, the first column refers to the relative order. For simplicity, it begins at one. The second column lists the time the action happened. The third column lists the dialogue between the tutor and the student. All the tutor turns start with “Tutor: ” while the student’s turns start with “Student: ”. In the last column, each of the student’s and tutor’s turns were mapped to the corresponding KCs. For example, in Line 2, the tutor made an ET decision and decided to tell. The target KC is KC_{20} . Because it is a tell, this turn is labelled with $*KC_{20}$ which means this is tutor-generated information in the dialogue on KC_{20} . Line 4 shows an example of an entry generated by the student, which is correct and also targeted to KC_{20} . So the turn is labelled with $+KC_{20}$.

In Table 5.1, there are total of five micro-step decisions. They are labelled as “D+ line number”, D2, D3, D5, D8, and D10 respectively. Among the five tutorial decisions, with the exception of D5, all remaining tutorial decisions involved KC_{20} . Therefore, to induce a KC-specific tutorial tactics on KC_{20} , only the tutorial decisions that involved KC_{20} were taken into account.

What follows is a simplified illustration of how the sample tutorial dialogue was transformed into one of eighteen features on KC_{20} . In Table 5.1 the feature I picked to illustrate is “tellsSinceElicitA”, one of three autonomy features. Initially the “tellsSinceElicitA”’s value is two, which means that the tutor made two tell decisions since the last elicit decision before this sample dialogue. For the decision D2, the tutor made another tell decision and now the “tellsSinceElicitA” value is three. The dialogue continues to the next decision, D3, and this time the tutor has made an elicit decision. So “tellsSinceElicitA”’s value goes back to 0. The next micro-step decision involving KC_{20} is D8. Because the tutor did not make any tell decisions between lines 3 and 8 the “tellsSinceElicitA” value is still 0. The tutor made two decisions in D8: first on JS and then on ET. The tutor decided not to skip the justification step and also decided to elicit the step from the student. So “tellsSinceElicitA” value remains 0 until the tutor made the decision in D10. At D10, the tutor made a decision to tell and thus “tellsSinceElicitA”’s value became 1 until the tutor was ready to make another

decision.

Table 5.1: A Simplified Example of Part of Student Log on Training Problem P4

#	Time	String	Label
1	14:56:08	Tutor:So let's start with determining the value of KE0.	
2	14:56:11	Tutor:To calculate the rock's kinetic energy at T0, let's apply the definition of kinetic energy.	<u>TELL</u> *KC₂₀, D2
3	14:56:27	Tutor>Please write the equation for applying the definition of kinetic energy to the rock at T0.	<u>ELICIT</u> KC₂₀, D3
4	14:56:31	Student:ke0 = 1/2 * M * V0^2	+KC ₂₀
5	14:56:32	Tutor:Since all the variables in the equation are known except for KE0, we can calculate KE0 easily. What is the value of KE0?	<u>ELICIT</u> KC₂₃, D5
6	14:56:47	Student:1.2J	+KC ₂₃
7	14:56:48	Tutor:Okay.Before we go any further, let's discuss this principle in a little more detail.	
8	14:56:56	Tutor:Why does the rock have a non-zero kinetic energy at T0?	<u>JUSTIFY ELICIT</u> KC₂₀, D8
9	14:57:17	Student: because it has a nonzero mass and a nonzero magnitude of velocity	+KC ₂₀
10	14:57:18	Tutor: Yes. We do not know the direction of the rock's velocity at T0 from the problem statement. However, the direction of v0 does not affect the rock's kinetic energy at T0.	<u>TELL</u> *KC₂₀, D10

Table 5.2 illustrates how the three autonomy features got updated as the sample dialogue in Table 5.1 continues. In Table 5.2, the first row shows the corresponding values for the three autonomy features when the tutor reaches D2 in Figure 5.1 and the last row shows their values when the tutor reaches the decision D10 in Figure 5.1. The rest of the 18 features were calculated in a similar manner. Appendix J show how the eighteen features varied as the sample example dialogue in Table 5.1 goes on.

Table 5.2: Autonomy Features Updated

Decision	tellsSinceElicitA	pctElicitA	pctTellsKCSessionA
D2	2	0.47	0.64
D3	3	0.47	0.65
D8	0	0.48	0.64
D10	0	0.49	0.63

5.1.4.2 Maximum Number of Features Previously I discussed the problems of data sparsity for RL. In an RL model, the size of the state space increases exponentially as the number of involved features increases. In order to learn effective tutoring tactics, a corpus should cover each of these states at least once, which means at least 2^{18} in our case. However, it is almost impossible to do so due to the high cost of collecting educational data. On the other hand, the learned policy may become too subtle to be necessary. Based on the size of the Exploratory Corpus collected in Study 1 and the number of categories of features defined in this study is four, the state representation was capped four features. Moreover, as discussed in subsection 5.1.4.1, the maximum number of features within each of the four categories was limited to one. This was done because it was anticipated that this would better represent the relevant information.

5.1.4.3 Feature Discretization In Study 2, except for EarlyTrainingPS(F7, binary), ProblemComplexityPS (F8, three), and DuringWalkThroughPS (F9, binary) which are dis-

crete features, the remaining fifteen features are continuous features and need to be discretized. In Study 2, this was accomplished for each feature by turning each feature into a binary feature via a median split. This was done in order to balance the number of cases across clusters. Therefore, apart from ProblemComplexity, which has three categories, all other features are binary.

5.1.4.4 Feature Selection In Study 2 a two-pass feature selection process was employed. For each KC and decision pair (e.g. KC_1 , ET) 18 single-feature MDPs were generated. Each MDP used one and only one of the 18 features to represent the state, and used the relevant tutorial decisions to represent the actions. All other features were ignored. For each of these MDPs, the Tetreault and Litman’s toolkit was applied [Tetreault and Litman, 2008] to induce a *single-feature policy* together with its corresponding ECR.

In the second pass the four best features were selected. Using Moore’s categories and the ECR, the policy with highest ECR from each of the categories was selected. The process involved choosing one from among features 1-3 for the single autonomy feature; one from among features 4-6 for the temporal situation feature, and so on. The criteria used for the selection ECR, specifically the single-policy feature who had the highest ECR relative to its’ peers is selected. Note that the policy’s Confidence Interval was ignored because ECR was more widely used in RL community for evaluating the derived policies [Williams and Young, 2007b, Williams and Young, 2007a, Janarthanam and Lemon, 2009]. An MDP was then defined for the KC and action decisions by using the four lead features for the state representation, and induced a new *four-feature policy* from it. From all eighteen single-feature policies and the four-feature policy, the policy with the highest ECR is selected for each KC and action decision. Recall that the higher the ECR of a policy, the more effective the policy is supposed to be.

However, a subsequent analysis of this feature selection method showed its limitations. Notably, other feature selection methods were applied to the 18 features, which included four RL-based feature selection methods (reviewed in Chapter 6) and a random feature selection method, the induced policies had significantly higher ECR [Chi et al., 2008a]. For example, Table 5.3 shows the ECR of the DichGain tutorial tactics and the tutorial tactics

induced by applying new feature selection approaches to the eight primary KCs for the ET decisions. The new tutorial tactics had higher ECRs than the DichGain ones across all eight KCs. For example, line 5 shows that $ECR(\pi_{DichGain}(KC_{22}, ET)) = 9.4$ while $ECR(\pi_{new}(KC_{22}, ET)) = 44.29$, four times higher than the former.

Table 5.3: Compare DichGain Tactics With Tutorial Tactics Under New Feature Selection Methods On Eight Primary KCs

		DichGain Tutorial Tactics			Tutorial Tactics With New Feature Selection		
		ECR	95% CI	Range	ECR	95% CI	Range
1	KC_1	20.196	[5.19, 34.19]	29.01	51.79	[32.67, 63.71]	31.04
2	KC_{14}	54.15	[47.9, 59.58]	11.69	59.48	[54.3, 63.21]	8.91
3	KC_{20}	4.81	[0.75, 8.66]	7.9	8.08	[4.24, 11.9]	7.66
4	KC_{21}	15.48	[7.85, 21.78]	13.94	26.94	[19.8, 29.28]	9.48
5	KC_{22}	9.4	[-5.37, 20.69]	26.05	44.29	[23.49, 50.51]	27.02
6	KC_{24}	7.23	[2.72, 11.31]	8.59	12.91	[7.22, 16.43]	9.21
7	KC_{27}	16.78	[5.95, 24.9]	18.95	27.25	[13.87, 32.16]	18.29
8	KC_{28}	15.29	[2.52, 26.05]	23.52	32.8	[22.08, 38.61]	16.53

5.1.5 Conflicting Policies

In some cases, a given tutorial step involves multiple KCs and thus, multiple policies. When multiple policies are relevant, the policy with the highest ECR was followed.

5.1.6 Summary: Procedure of Inducing Tutorial Tactics in Study 2

In sum, Study 2 involved using the Exploratory Corpus collected in Study 1 as the training corpus and the reward functions are defined as either +100 (high learner) or -100 (low learner) based on corresponding KC-based NLGs. In Study 2 a total of thirty-three policies

were induced: twenty-two ET policies (21 KC-specific and one KC-general) and eleven JS policies (10 KC-specific and one KC-general).

In order to induce the necessary policies a multi-pass approach was adopted. In the multi-pass approach, a set of 18 MDPs were constructed for each KC and decision pair (e.g. KC4, JS) with all features either discrete or discretized via a median split. Then a set of 18 single-feature policies was induced, one for each MDP. For each of Moore’s four categories, a feature was selected whose corresponding single-feature policy had the highest ECR among features in the same category. A four-feature MDP was then defined from which a more complex policy was induced. All of the work here was done with Tetreault and Litman’s toolkit [Tetreault and Litman, 2008]. Finally, the policy with highest ECR for the KC and decision pair was selected from among the 19 derived policies: 18 single-feature policies and one four-feature policy.

Appendix K lists the 22 ET tutorial tactics and the 11 JS policies that were induced and applied in Study 2. One of the resulting policies is shown in Figure 5.1. This policy involved four features: durationKCT, ProblemComplexityPS, tellsSinceElicitA, and pctCor-

1. *Features:* durationKCT, ProblemComplexityPS, tellsSinceElicitA, pctCorrectKCSessionPM
2. *Cutoff:* durationKCA = '50.0' tellsSinceElicitA = '0.0001' pctCorrectKCSessionPM = '0.7179'
3. *Policy:*
 - a. *Elicit:* **0:MED:1:0**, 1:COMP:1:0, 0:COMP:1:1, 0:MED:0:0, 0:COMP:1:0, 0:MED:1:1, 0:COMP:0:1, 1:COMP:0:1
 - b. *Tell:* 1:MED:0:1, 1:MED:0:0, 1:MED:1:0, 1:MED:1:1, 0:MED:0:1
 - c. *Else:* 0:COMP:0:0, 1:COMP:0:0, 1:COMP:1:1

Figure 5.1: The Induced Policy $\pi_{DichGain}(KC_{21}, ET)$: Gravitational Potential Energy

rectKCSessionPM. Three of these features were continuous. The median cutoff values used to discretize them are shown in line 2 (“cutoff”). This policy contains a total of $2^4 = 16$ rules. In eight cases, the tutor elicited (line a); in five cases the tutor elected to tell (line b); and in the remaining three the tutor could choose to do either (line c). For example, “elicit: [0:MED:1:0]” (shaded in line a) means:

IF the duration since the most recent decision made on KC_{21} is less than 50sec;
AND the ProblemComplexity is 'medium';
AND the students have received at least one tell since the most recent elicit (tellsSinceElicit)
AND the student’s performance on this kc in today’s session is less than 71.79% correct;
THEN: the tutor should elicit the next step from the student.

The example in Figure 5.1 indicated that the induced tactics were a very specific set of case decisions, and could easily be implemented back into Cordillera. Moreover, the tactics were quite subtle.

Then the thirty-three induced policies were implemented back into Cordillera producing a version of the system called DichGain-Cordillera. This version of the system used the KC-specific policies when facing a relevant decision, resolved ties by selecting the policy with the best ECR, or followed the KC-general strategies when no policy was relevant. As described below, a set of students were trained on this system to collect the DichGain Corpus. For Study 2, the author was responsible for application of the MDP toolkit to the Exploratory Corpus and induced the 33 DichGain tutorial policies. Once the policies were implemented back into Cordillera, the author acted as a human wizard during the collection of the DichGain Corpus.

5.2 METHODS

5.2.1 Participants

Data was collected over a period of three months during Spring 2008. As in Study 1, a set of forty-two college students were recruited and paid for their time regardless of completion.

The students were required to have a basic knowledge of high-school algebra and not to have enrolled in college-level physics courses. All told, thirty-seven students completed the experiment.

5.2.2 Materials & Procedures

The students followed the same procedure, used the same preparatory materials and problems, and involved the same group of human wizards as in Study 1. More specifically, the DichGain group completed a background survey, read a textbook covering the target domain knowledge, took a pre-test, solved the same seven training problems in the same order on DichGain-Cordillera, and finally took a post-test. Only two salient differences exist between Study 1 and Study 2:

1. Interaction decisions made by DichGain-Cordillera were guided by thirty-three derived tutorial tactics; and
2. One test problem, Q_{20} , on the pre- and post-test was changed for Study 2 to, Q_{20}^* . Both Q_{20} and Q_{20}^* are multiple-choice questions and cover the same KCs. But Q_{20} is a simple question and had only two choices (true, false), so there is a good chance that students could guess the answer. The new version of Q_{20}^* , covered the same KCs but was more difficult by providing five choices. So it is less likely that students could guess the answer. The remaining 32 test items were identical in both studies. And as with Study 1, the pre- and post-tests in Study 2 were identical.

5.2.3 Grading

All tests were graded by the same grader as in Study 1 (not the author). She applied the same grading metrics and carried out the same grading process resulting in both the overall and KC-based grades.

5.2.4 Measures

There were two research objectives in Study 2: first, to determine whether the DichGain group learned by training on DichGain-Cordillera, and second, to examine the DichGain group and the Exploratory group to see whether the induced DichGain tutorial tactics would result in better learning performance than making random decisions. Note that given the cost of the study, we did not run a strict control-experimental study but simply confirmed whether the trend of the DichGain group over-performing the Exploratory group existed. The two groups' learning performances were compared, using both the students' pre-test, post-test, adjusted post-test scores and NLG, under both the overall grading criteria and the cumulative KC-based grading criteria. The adjusted post-test can be measured as a linear association between the real post-test score for each student and the difference between the pre-test score for the students and the mean of pre-test scores. The formula for the adjusted post-test score is:

$$posttest_i^* = posttest_i - \beta \times (pretest_i - \overline{pretest}) \quad (5.1)$$

where i stands for the student, $posttest_i^*$ for the adjusted post-test score for student i , $posttest_i$ for the true post-test score for the student i , β is the regression coefficient of the post-test score upon the pre-test score, $pretest_i$ is the true pre-test score for the student i , and $\overline{pretest}$ is the mean of the pre-test scores.

Results showed that there was no significant difference between the DichGain and Exploratory groups under either grading criteria. There are two potential reasons for this. One is the lack of random assignments and two is that the RL approach may be limited. As a result, Study 3 focused on a full-scale comparison by exploring a wider range of methods to deal with the five RL issues. For example, three training corpora were explored in Study 3: the Exploratory Corpus collected in Study 1, the new DichGain Corpus in this study, and a combination of the two in a new corpus. Because of this decision, the second part of the results section will focus on the characteristics of the DichGain corpus. These characteristics will include the general number of decisions the tutor made, the number of ET decisions and

the I-ratio, and the number of JS decisions and the J-ratio. Additionally, because the DichGain corpus will be used to induce KC-specific tutorial tactics in Study 3, the description will include some KC-based learning performance and corpus characteristics as well.

5.3 RESULTS

In Study 2, it took each student from three to six sessions to complete the study. These sessions were spaced over a period of one to three weeks. The sessions generally took less than two hours to complete. The students spent roughly five hours, ranging from as few as four hours to as many as nine hours, training on DichGain-Cordillera.

5.3.1 Compare Pre- and Post-test

Table 5.4: DichGain Students' Pre- vs. Post-test Performance

		Min	Max	Mean	σ
Overall Grading	Pretest	0.04	0.74	0.40	0.18
	Posttest	0.18	0.96	.58	.19
	NLG	-0.09	0.89	0.33	.21
Cumulative KC-based Grading	Pretest	0.04	0.77	.42	.17
	Posttest	0.08	0.97	.54	.20
	NLG	-0.33	0.86	0.25	0.23

A one-way ANOVA was used to test for performance preference differences between the pre- and posttests. Regardless of grading criteria, participants made significant gains from pre-test to post-test (Table 5.4), $F(1, 72) = 16.86, p = .000, R^2 = .69$ under the overall grading criteria and $F(1, 72) = 8.55, p = .005, R^2 = .71$ under the cumulative KC-based grading criteria. The overall NLG scores ranged from -0.09 to 0.89 ($M = 0.33, SD = .21$).

The cumulative KC-based NLG scores ranged from -0.33 to 0.86 ($M = 0.25, SD = .23$). Table 5.4 summarize the minimum, maximum, mean, and SD values for each test scores.

Figure 5.2 shows the DichGain group’s pre- and post-test scores under the two grading criteria. A double asterisk (**) indicates that the difference is statistically significant ($p < 0.05$). To summarize, the DichGain group scored significantly higher in the post-test than in the pre-test. .

5.3.2 Post-hoc Comparison: DichGain vs. Exploratory

5.3.2.1 Post-hoc Comparison: DichGain vs. Exploratory On Training Time In a post-hoc comparison with a one-tailed paired t-test, there were no significant overall time on task differences between the DichGain group ($M = 294.33, SD = 87.50$) and the Exploratory group ($M = 280.38, SD = 66.88$) across the seven training problems : $t(99) = .88, p = .38$.

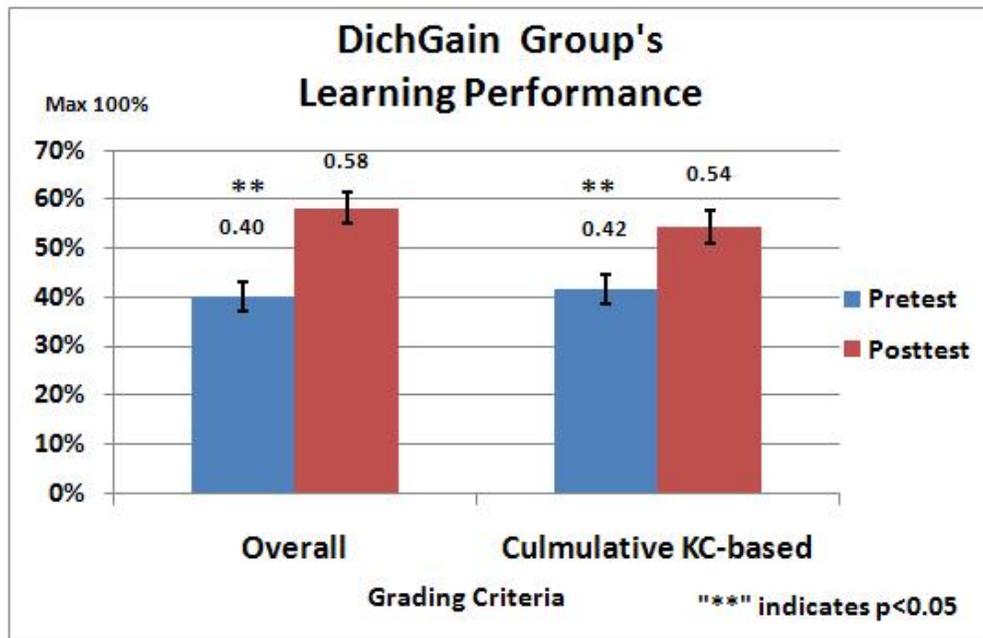


Figure 5.2: Learning Performance of Exploratory Group

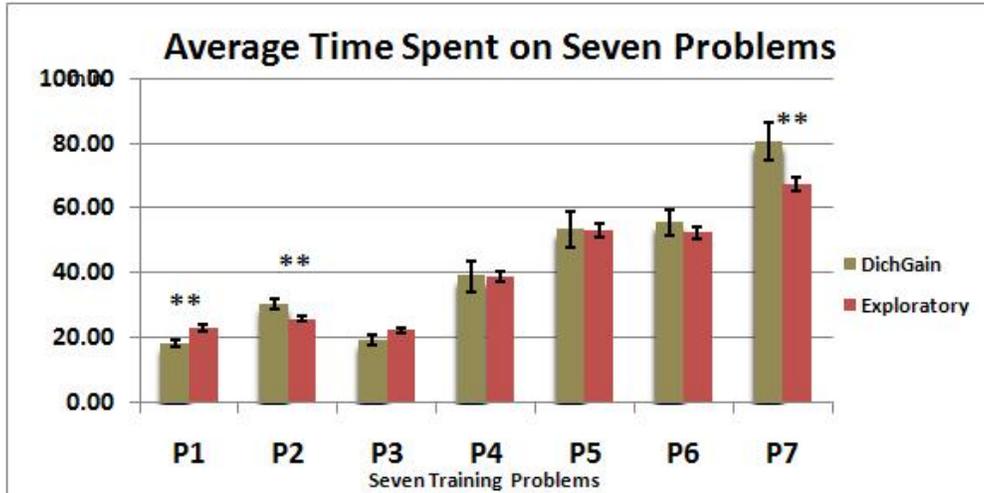


Figure 5.3: Per Problem Time Comparison: DichGain vs. Exploratory Group

However, a significant difference was found between the two groups in the time they spent on P1, P2 and P7 Figure 5.3 compared the average time students spent on each training problem between the two groups. On P1 the DichGain group spent significantly less time than the Exploratory group with a one-tailed paired t-test ($t(98) = 3.15, p = .002$) while on P2, the DichGain group spent longer than the Exploratory group ($t(99) = 2.56, p = .012$). Similarly on P7, the DichGain group spent significantly longer time than the Exploratory group: ($t(99) = 2.46, p = .016$).

5.3.2.2 Post-hoc Comparison: DichGain vs. Exploratory On Learning Performance

Because of an administrative error, all of the background information for DichGain group was not available for comparison. As mentioned above, one test problem Q20 was changed from Study 1 to Q20* for Studies 2 and 3. So in order to compare the two groups, Q20 and Q20* were excluded from the scores used here. As described in the previous chapter, the tests contained thirty-three test items which covered 168 KC occurrences. Removing Q20 or Q20* reduced this total by one leaving thirty-two test items covering 166 KC occur-

Table 5.5: DichGain vs. Exploratory Scores: Pre vs. Post-Test (No Q_{20})

		Dich $m(\sigma)^a$	Exp $m(\sigma)^b$	Stat ^c	d	$1 - \beta$
Overall	Pre-	.40 (0.18)	.47 (0.20)	$t(99) = 1.86, p = 0.066$	-0.37	0.39
	Post-	0.58 (0.19)	0.61 (0.18)	$t(99) = 0.78, p = 0.44$	-0.16	0.54
	Adj. Post-	0.62 (0.10)	0.59 (0.10)	$F(1, 98) = 1.99, p = 0.16$	0.3	0.46
	NLG	0.34 (0.20)	.28 (0.21)	$t = 1.36, p = 0.18$	0.29	0.47
Cumulative KC-Based	Pre-	0.41 (0.17)	0.45 (0.20)	$t(99) = 0.99, p = 0.32$	-0.21	0.5
	Post-	0.54 (0.20)	0.57 (0.21)	$t(99) = 0.66, p = 0.51$	-0.15	0.59
	Adj. Post-	0.57 (0.12)	0.56 (0.12)	$F(1, 98) = 0.09, p = 0.77$	0.08	0.78
	NLG	0.26 (0.23)	0.25 (0.26)	$t(99) = 0.23, p = 0.82$	0.04	0.82

^aThe Mean and SD of DichGain Group.

^bThe Mean and SD of Exploratory Group.

^cExcept an ANCOVA using pre-test score as the covariate on Adj.Post-test scores, the two groups were compared with one-tailed paired t-tests on the pre-test, post-test and NLG scores.

rences. In the subsections learning performance will be compared across both groups using both the overall and cumulative KC-based scores. For the overall scores, the maximum raw score was 32 points while for the cumulative KC-based score had a maximum of 166 points. For comparison purposes both scores were normalized to 1.

A one-way ANOVA was used to test for performance preference differences between the pre- and posttests across the two groups. Across 32 test questions, participants in the first two studies made significant gains from pre-test to post-test, $F(1, 200) = 35.88, p = .000, R^2 = .70$ under the overall grading criteria and $F(1, 200) = 19.51, p = .000, R^2 = 0.69$ under the cumulative KC-based grading criteria. In a post-hoc comparison, however, no significant pre-test score differences were found between the two groups on pre-test scores, post-test scores, adjusted post-test scores, and NLG under either the overall-grading rubric or the cumulative KC-based scores (Table 5.5). The first column in Table 5.5 shows the eight comparisons: pre-test scores, posttest scores, adjusted posttest scores, and NLG under both the overall-grading rubric and the cumulative KC-based scores. The second column in Table 5.5 lists the means (m) and SDs σ of two groups' corresponding scores. The third column lists the corresponding statistical comparisons. No significant difference was found

between the two groups across the eight comparisons. However, the DichGain students did demonstrate marginally significant lower pre-test scores than the Exploratory group under the overall grading criteria only.

The fourth column lists the effect size of the comparison. There are several accepted ways to measure effect size, such as Cohen's d effect sizes based on means, Hedges' g and so on. For this dissertation, Cohen's d is selected and it is defined as the mean learning gain of the experimental group minus the mean learning gain of the control group, divided by the groups' pooled standard deviation. The final column listed the statistical power of the comparison, $1 - \beta$ ¹. Generally speaking, it must be kept correspondingly high. Ideally, power should be at least 0.80 to detect a reasonable departure from the null hypothesis. The reward functions used for inducing DichGain tutorial tactics were based on the students' cumulative KC-based NLGs or KC-based NLG scores. However, the last row in Table 5.5 shows that the Exploratory and DichGain groups were not significantly different on cumulative KC-based NLGs and its power reached an acceptable level: 0.82 (often considered to be between .80 and .90).

Although no significant difference was found between two groups on learning performance and overall time on training, the DichGain students did have a marginally significant lower pre-test score than the Exploratory group under the overall grading criteria. One potential reason for an absence of difference in learning between the two groups may be because the lack of random assignment. However, there are other potential reasons for this. For example, it might be because of the limitation of the RL approach used in Study 2. As discussed above, the feature selection method in Study 2 is somewhat greedy-like. So in Study 3, significantly more feature selection methods were explored to find ways to better use RL and a full-scale comparison was run.

At this point in the research, two training corpora existed: the Exploratory Corpus in which all decisions were randomly made, and the corpus that was collected by following the induced DichGain tutorial tactics induced from the Exploratory Corpus. Although there was no significant learning performance difference between the two groups, the two corpora

¹ β represents Type II error: false negative. It refers to the error of failing to reject a null hypothesis when it is in fact not true.

may differ in other aspects. In Study 3, both the Exploratory and the DichGain corpora were used as training corpora. The characteristics of the Exploratory Corpus were discussed previously in Chapter 4 and the characteristics of the DichGain corpus will be discussed below.

Finally, the two corpora will be compared using the following measurements: the average number of ET decisions, JS decisions, and overall decisions that the tutor made. Another point of comparison is measured by the I-ratio, the J-ratio, and the number of justification steps. In addition, because the DichGain corpus will be used to induce KC-specific tutorial tactics in Study 3, its KC-based learning performance and corpus characteristics will also be discussed.

5.3.3 Post-hoc Comparison: DichGain vs. Exploratory Tutorial Corpora

The DichGain corpus was used as one of the training corpora to derive KC-based tutorial tactics in Study 3. The decision to choose one corpus over another means that it is valuable to compare the characteristics of the two corpora. Similar to the Exploratory corpus, each student's individual problem dialogues were combined into a single super-dialogue listing all tutor-student interactions in order of occurrence. Thus, one tutorial dialogue was combined per participant.

5.3.3.1 Post-hoc Comparison: DichGain vs. Exploratory On Overall Tutorial Decisions Table 5.6 compares the various tutorial decisions with a one-tailed paired t-test across all KCs between the DichGain and Exploratory Corpora. Except for the total number of overall decisions and the total number of ET decisions, the two corpora differed on all the other seven aspects (labeled with "***"). Overall, the DichGain Corpus is significantly less interactive in that the DichGain students received more tells and less elicits from the tutor than the Exploratory Corpus. As a result, the I-ratio of the DichGain corpus was significantly lower than that of the Exploratory corpus. Moreover, the DichGain-Cordillera skipped more and executed less justification steps than the Exploratory-Cordillera.

Table 5.6: Overall Tutorial Decision Characteristics: DichGain vs. Exploratory Corpora

Decision	Condition	Mean	σ	Stats
tell**	DichGain (37)	152.46	13.05	$t(99) = 6.663, p = 0.000$
	Exploratory (64)	138.02	8.71	
elicit**	DichGain (37)	118.08	13.33	$t(99) = -6.956, p = 0.000$
	Exploratory (64)	135.88	11.82	
ET decisions	DichGain (37)	270.54	10.00	$t(99) = -1.396, p = 0.166$
	Exploratory (64)	273.89	12.46	
skip-Justify**	DichGain (37)	33.54	4.80	$t(99) = 7.728, p = 0.000$
	Exploratory (64)	26.44	4.24	
Justify**	DichGain (37)	24.89	3.59	$t(99) = -6.826, p = 0.000$
	Exploratory (64)	30.17	3.83	
JS decisions**	DichGain (37)	58.43	2.81	$t(99) = 2.742, p = 0.007$
	Exploratory (64)	56.61	3.43	
Overall Decisions	DichGain (37)	307.57	12.45	$t(99) = 0.749, p = 0.456$
	Exploratory (64)	305.48	14.01	
I-ratio**	DichGain (37)	0.44	0.04	$t(99) = -7.967, p = 0.000$
	Exploratory (64)	0.50	0.03	
J-ratio**	DichGain (37)	0.43	0.07	$t(99) = -7.894, p = 0.000$
	Exploratory (64)	0.53	0.06	

5.3.3.2 Post-hoc Comparison: DichGain vs. Exploratory On Individual KCs

Table 5.7 shows the number of tutorial decision steps for each KC and each type of tutorial decision in DichGain Corpus. The third and fourth columns list the number of ET and JS tutorial decisions for the KC. The last column lists the statistical results of comparing the KC-based pre-test scores with post-test scores with one-tailed paired t-tests. If the DichGain group had significantly higher post-test scores than its pre-test scores, the corresponding KC

were labelled with “**” next to their name in (column 2).

From Table 5.7, it can be seen that, as with the Exploratory Corpus, the DichGain corpus had twenty-one KCs that occurred in at least one ET tutorial decision step and ten KCs for JS decisions. Additionally, the number of occurrences of ET decisions varied from one to seventy-two occurrences; for JS decisions it varied from two to sixteen occurrences. Among the eight primary KCs, students learned significantly from six of them: KC_1 , KC_{14} , KC_{20} , KC_{21} , KC_{22} , KC_{24} but not on KC_{27} and KC_{28} .

Table 5.7: Tutorial Decisions Per KC.

	KC	ET	JS	Total	Pre- and Posttests
1	KC_1^{**}	4.16		4.16	$t(72) = 2.80, p = 0.01$
2	KC_3	2.11		2.11	$t(72) = 0.66, p = 0.51$
3	KC_5	7.05		7.05	$t(72) = 1.64, p = 0.11$
4	KC_9^{**}	1.22		2	$t(72) = 2.46, p = 0.02$
5	KC_{12}	1.46	2	3	$t(72) = 1.27, p = 0.21$
6	KC_{13}	3		3	$t(72) = 1.27, p = 0.21$
7	KC_{14}^{**}	8.38	2	9	$t(72) = 2.03, p = 0.05$
8	KC_{15}	7.62		8.11	$t(72) = 0.98, p = 0.33$
9	KC_{17}	4.43		5.43	$t(72) = 1.68, p = 0.10$
10	KC_{18}	1.14		2	$t(72) = 1.60, p = 0.11$
11	KC_{20}^{**}	72.43	16.43	82.84	$t(72) = 4.45, p = 0.00$
12	KC_{21}^{**}	49	16.86	59.05	$t(72) = 4.30, p = 0.00$
13	KC_{22}^{**}	31.32	3.97	32.78	$t(72) = 2.41, p = 0.02$
14	KC_{23}^{**}	61.11	4	62.95	$t(72) = 3.39, p = 0.00$
15	KC_{24}^{**}	48.97	15.89	60.24	$t(72) = 3.11, p = 0.00$
16	KC_{25}^{**}	9.43		9.78	$t(72) = 2.65, p = 0.01$
17	KC_{26}	5.41	2.81	7.46	$t(72) = 1.76, p = 0.08$
18	KC_{27}	21.92	4.62	23.73	$t(72) = 1.99, p = 0.051$
19	KC_{28}	14.19	5.14	18.32	$t(72) = 1.82, p = 0.07$
20	KC_{31}	18.32		19.16	$t(72) = 1.82, p = 0.07$
21	KC_{32}^{**}	14.24		14.24	$t(72) = 2.29, p = 0.03$

Overall, there was a significant difference between the two corpora on the I-ratio and J-ratio. However, as this difference is analyzed, the variance becomes more complex. Table 5.8 shows the I-ratio difference between the DichGain and Exploratory corpora with one-tailed

paired t-tests. The third and fourth columns of the table list the mean of the I-ratio in the DichGain Corpus and Exploratory Corpus respectively. The fifth column gives the direction of the difference in which “ $DG > EX$ ” represents that the DichGain Corpus was more interactive than the Exploratory Corpus on corresponding KCs. Similarly, “ $DG < EX$ ” means the reverse is true. If the column is blank, it means that there were no significant differences between the corpora on the I-ratio for the corresponding KC. The last column shows the statistical results between the two corpora. If the difference is significant, the KC name in column 2 is labeled with “**”.

As shown in Table 5.8 the corpora differed significantly in terms of I-ratio on all but three KCs: KC_9 , KC_{22} , and KC_{32} . The DichGain group was significantly less interactive than the Exploratory group on six KCs (KC_{12} , KC_{18} , KC_{20} , KC_{23} , KC_{24} , and KC_{27}), and significantly more interactive than the Exploratory group on the remaining twelve KCs.

Table 5.8: I-Ratio Between DichGain vs. Exploratory on a per-KC basis.

	KC	DichGain	Exploratory	Diff	Stats Comparison
1	KC_1^{**}	0.90	0.49	$DG^a > EXP^b$	$t(99) = 8.92, p = 0.0000$
2	KC_3^{**}	0.80	0.51	$DG > EXP$	$t(99) = 4.05, p = 0.0000$
3	KC_5^{**}	0.56	0.46	$DG > EXP$	$t(99) = 3.37, p = 0.0010$
4	KC_9	0.35	0.47		$t(99) = -1.35, p = 0.1810$
5	KC_{12}^{**}	0.21	0.52	$DG < EXP$	$t(99) = -3.81, p = 0.0000$
6	KC_{13}^{**}	0.90	0.45	$DG > EXP$	$t(99) = 8.17, p = 0.0000$
7	KC_{14}^{**}	0.89	0.50	$DG > EXP$	$t(99) = 12.00, p = 0.0000$
8	KC_{15}^{**}	0.60	0.47	$DG > EXP$	$t(99) = 4.38, p = 0.0000$
9	KC_{17}^{**}	0.71	0.48	$DG > EXP$	$t(98)4.66, p = 0.0000$
10	KC_{18}^{**}	0.04	0.43	$DG < EXP$	$t(99) = -5.39, p = 0.0000$
11	KC_{20}^{**}	0.23	0.50	$DG < EXP$	$t(99) = -17.26, p = 0.0000$
12	KC_{21}^{**}	0.65	0.50	$DG > EXP$	$t(99) = 6.76, p = 0.0000$
13	KC_{22}	0.46	0.49		$t(99) = -1.53, p = 0.1300$
14	KC_{23}^{**}	0.45	0.49	$DG < EXP$	$t(99) = -3.36, p = 0.0010$
15	KC_{24}^{**}	0.28	0.50	$DG < EXP$	$t(99) = -17.43, p = 0.0000$
16	KC_{25}^{**}	0.73	0.50	$DG > EXP$	$t(99) = 7.70, p = 0.0000$
17	KC_{26}^{**}	0.58	0.43	$DG > EXP$	$t(99) = 3.05, p = 0.0030$
18	KC_{27}^{**}	0.41	0.51	$DG < EXP$	$t(99) = -3.78, p = 0.0000$
19	KC_{28}^{**}	0.64	0.48	$DG > EXP$	$t(99) = 4.23, p = 0.0000$
20	KC_{31}^{**}	0.67	0.51	$DG > EXP$	$t(99) = 7.42, p = 0.0000$
21	KC_{32}	0.52	0.50		$t(99) = 0.55, p = 0.5830$

^aDG = DithGain

^bEXP=Exploratory

Table 5.9 shows the J-ratio difference between the two corpora. Similarly, the third and fourth columns of the table list the mean of the J-ratio in the DichGain Corpus and Ex-

Table 5.9: Justify Ratio Differences on a per-KC Basis.

	KC	DichGain	Exploratory	Diff	Stats Comparison
1	KC_{12}^{**}	0.23	0.54	$DG < EXP$	$t(99) = -4.40, p = 0.0000$
2	KC_{14}^{**}	0.69	0.48	$DG > EXP$	$t(99) = 3.44, p = 0.0009$
3	KC_{20}^{**}	0.42	0.51	$DG < EXP$	$t(99) = -4.24, p = 0.0001$
4	KC_{21}^{**}	0.46	0.56	$DG < EXP$	$t(99) = -4.27, p = 0.0000$
5	KC_{22}	0.63	0.61		$t(99) = 0.25, p = 0.8043$
6	KC_{23}^{**}	0.54	0.78	$DG < EXP$	$t(99) = -6.72, p = 0.0000$
7	KC_{24}^{**}	0.33	0.57	$DG < EXP$	$t(99) = -8.58, p = 0.0000$
8	KC_{26}^{**}	0.84	0.60	$DG > EXP$	$t(99) = 3.68, p = 0.0004$
9	KC_{27}^{**}	0.63	0.51	$DG > EXP$	$t(99) = 2.49, p = 0.0143$
10	KC_{28}^{**}	0.36	0.54	$DG < EXP$	$t(99) = -3.99, p = 0.0001$

ploratory Corpus respectively. The fifth column gives the direction of the difference in which “ $DG > EX$ ” demonstrates that DichGain Corpus got justification steps more frequently than the Exploratory Corpus on corresponding KC, while “ $DG < EX$ ” means the reverse is true. If this column is blank, it means that there were no significant differences between the corpora on the justification ratio on the KC. The last column shows the statistical results between the two corpora with one-tailed paired t-tests. If the difference is significant, the KC name in column 2 was labeled with “***”. In Table 5.9. There was no significant difference in terms of the number of justifications between the two corpora on KC_{22} . While on six (KC_{12} , KC_{20} , KC_{21} , KC_{23} , KC_{24} , and KC_{28}) the DichGain group were more likely to skip a justification step than the Exploratory group. There were only three instances for two KCs (row: KC_{14} , KC_{26} , and KC_{27}) in which the DichGain group was more likely to receive tells than the Exploratory group.

Thus, although no significant learning differences were found between the two groups for time on task or learning performance, significant differences were found between the two corpora. The DichGain corpus was significantly less interactive and included less justification

steps than the exploratory corpus.

5.4 DISCUSSION

The goals in Study 2 were to investigate how to apply RL to induce tutorial tactics from a training corpus and then to test whether the induced tutorial tactics would result in more effective learning performance than making random decisions. This was to be accomplished without running a full-scale comparison. Results showed that following the DichGain tutorial tactics generated significantly less elicits and included fewer justification steps than following the random decisions in the Exploratory group. A more detailed analysis, however, showed that this difference varied from KC to KC. While applying RL did induce tutorial tactics from the Exploratory corpus and the induced tutorial tactics were subtle, they did not seem to be more effective. Despite of the lack of random assignment, no significant difference was found between the two groups on either the pre-test, post-test, adjusted post-test or the NLG.

There were at least three potential reasons for lack of difference in learning performance between the DichGain and Exploratory groups. First, a full comparison of the DichGain and Exploratory groups was not run by assigning students randomly into the two groups. Second, the hypothesis may simply be incorrect, that micro-level policies covering interactive decisions like ET and JS do not affect students' learning. The decisions may be too "fine-grained" to have a real impact on learning, no matter how optimal the policy. Initial analysis based on the comparison of the DichGain and Exploratory groups appears to support previous research. That research suggests that given that content is controlled to be same, pedagogical tutorial tactics may not result in different learning. Third, it is also possible that lack of a difference in learning performance may be caused by limitations in the RL approach.

In other words, applying RL to induce tutorial tactics may not be a simple task for which we can plug a toolkit into the training corpus and induce effective tutorial tactics. As demonstrated in Study 2, tutorial tactics depend on many factors, such as feature choices,

feature selection, feature discretization and so on. Their effectiveness might also depend on how we implement them back into Cordillera. For example, how we deal with conflicting policies. It can be argued that in Study 2, exploration of these factors was limited. For example, only eighteen features were included in our search space, but no more than four appeared in the final induced tutorial tactics. It is possible that the selected features were insufficient to adequately represent the state space. Moreover our greedy-like feature selection process and the discretization procedure of using simple median splits may also have limited our success.

Study 3 was designed to address these reasons in hopes of producing more effective pedagogical tutorial tactics. In it the approach to RL-related issues was modified. For example, the training dataset was expanded to include both the Exploratory Corpus and the DichGain Corpus in the induction process. Also, more features were included in the feature states. To address the more weighty issue of learning performance, one set of tutorial tactics, like the policies in the present study, was derived with the goal of enhancing the tutorial decisions that contribute to the students' learning; while the other was derived with the goal of enhancing those decisions that contribute less or even none to the students' learning. To summarize, in contrast to Study 2, Study 3 included multiple datasets, a larger feature set, induction of policies based on multiple corpora, and random assignment of subjects to two comparable groups. The methods and outcomes of Study 3 are discussed in the next two chapters.

6.0 APPLYING RL TO INDUCE NORMALIZED GAIN (NORMGAIN) AND INVERSE NORMALIZED GAIN (INVNORMGAIN) TUTORING TACTICS

The conclusion of Chapter 5 identified three potential problems with Study 2 that might explain the absence of a learning difference in the two groups' performance. The earlier study did not run a full comparison by randomly assigning students into the two groups. The feature space selection may have been inadequate. The lack of a learning difference may also suggest that decisions on the level of elicit/tell (ET) and justify/skip-justify (JS), however well timed, cannot significantly affect the students' performance. Many previous studies showed that after solving the same training problems with the tutorial scripts written by the same authors, no significant difference was found among students' learning by means of different learning treatments [VanLehn et al., 2007a]. In this study, the content was controlled to be equivalent even at a much lower level than in these previous studies. Therefore, it is possible that these micro-decisions would not make a difference in students' learning. As shown in Studies 1 and 2, both the Exploratory and DichGain groups gained significantly. However, no significant difference was found between the two groups in a post-hoc comparison.

On the other hand, even if there was an impact by tutorial decisions on learning, random selection might have a good chance (50% chance given that both decisions were binary) to guess the "proper" decisions, and thus might have made enough effective decisions. If so, the impact of the tutorial decisions on learning would be canceled out. Therefore, in order to investigate whether micro-step decisions would make a difference in learning, the contrast between the two conditions in Study 3 was sharpened. Instead of choosing "random" as the control condition, the InvNormGain Group was selected.

In short, in Study 3 two sets of tutorial policies were induced: the Normalized Gain (NormGain) set induced by using the *students' NLG* as rewards and the Inverse Normalized Gain (InvNormGain) set was induced by specifically using *students' (1-NLG)* as rewards. In other words, the NormGain tutorial tactics were derived with the goal of enhancing the tutorial decisions that contribute to the students' learning, while the InvNormGain tutorial tactics were derived with the goal of enhancing those decisions that contribute less, or not at all, to the students' learning. If RL did live up to its promise, then it is expected that the NormGain students would out-perform their InvNormGain peers. This would occur if the micro-level decisions on ET and JS do impact learning.

Apart from the reward functions, the tactics were induced using the same general procedure. In this chapter, the main focus is on describing how RL was applied to induce these two sets of tutorial tactics. The experimental comparison of these two sets will be presented in Chapter 7. While the previous two studies were implemented by the ITR research group, Study 3 was designed, executed and evaluated by the author.

In order to induce tutorial tactics in Study 3, the same general learning procedure described in Chapter 3 and again in Chapter 5 was employed. As in the proceeding chapter the five major RL issues are addressed in Study 3 and changes made from Study 2 are explained.

6.1 TRAINING CORPUS

In Study 2, the only corpus available was the Exploratory Corpus. At this stage, three training corpora were available: the Exploratory Corpus collected in Study 1, the DichGain corpus from Study 2, and a combined corpus from both sets. The Exploratory Corpus consisted of 64 students' super-dialogues, the DichGain-Corpus consisted of 37 super-dialogues. The combined set contained 101. Each super-dialogue covered one student's entire interaction with the Cordillera system including all seven training problems.

The choice of Training Corpus is a complex one. As explained previously, the Exploratory Corpus was collected for RL and designed to explore the feature space evenly and without bias. The DichGain Corpus, by contrast, is similar to many other pre-existing corpora.

Inducing a successful policy from it would show the potential for applying RL to induce effective tutorial policies from most pre-existing data. The combined corpus, in theory, offers the benefits of both as well as an increased dataset. In this study, rather than selecting one corpus a priori, all three were used. More specifically, tutorial tactics were derived from each corpus separately, and then the best policies from all the sets were selected by ECR.

6.2 KNOWLEDGE COMPONENTS

Study 2 opted to induce tutorial policies for as many KCs as possible, covering every KC that was involved in at least one tutorial step. However, these KCs were not equally important. For example, in a domain such as physics, the domain principles are more challenging and important than other KCs. Consider, for example, KCs 23 and 20:

KC₂₃: The unit for energy is the Joule (J).

KC₂₀: If an object is moving, then its kinetic energy at a time is $\frac{1}{2}mv^2$, where m is the object's mass and v is the magnitude of the object's instantaneous velocity.

In domains such as physics, solving a problem requires producing an argument, proof or derivation consisting of one or more inference steps; each step is the result of applying a domain principle, operator or rule. Here *KC₂₀* is one of the major domain principles, i.e. the definition of Kinetic Energy, while *KC₂₃* is not a major principle. Therefore, *KC₂₀* is more important than the latter in that the student's overall learning performance depends more on learning a domain principle such as *KC₂₀* and less so on *KC₂₃*. Additionally, clearly *KC₂₀* is a complex principle with a non-trivial cognitive load while *KC₂₃* is an atomic fact and thus much simpler to convey and apply.

In Study 2 the ECRs of the KC-based tutorial tactics for *KC₂₃* are 42.45 on ET decisions and 47.22 on JS decisions. Either ECR is much higher than the corresponding ECR of the tutorial tactics on *KC₂₀*: 4.81 on ET decisions and 4.29 on JS decisions respectively. So when *KC₂₃* and *KC₂₀* co-occurred in a tutorial decision step, the dialogue manager would follow the policy for *KC₂₃* even though *KC₂₀* is a domain principle and learning it is more

important for students to learn the domain. Therefore, in Study 3, the decision was made to focus only on the eight primary KCs: KC_1 , KC_{14} , KC_{20} , KC_{21} , KC_{22} , KC_{24} , KC_{27} & KC_{28} , each of which represent a major domain principle shown in Table 4.1.

Table 6.1 compares the frequency and ratio of various tutorial decisions on the eight main KCs among the three corpora; the last row presents the comparison over *all* KCs across the three corpora. Columns 3 and 4 list the average number of ET and JS decisions per KC in each corpus. Column 5 shows the average number of tutorial overall decisions (regardless of whether it is ET or JS). Columns 6 and 7 present the I-ratio and J-ratio respectively. The last column presents a t-test comparison of the students' KC-based pre- and post-test scores. In the last column, if students' KC-based post-test scores were significantly greater than their corresponding pre-test scores, then the results of the t-test were listed in the last column. There were no cases in which students' post-test scores on a KC were significantly lower than their corresponding pre-test scores.

From Table 6.1, it can be seen that the average number of tutorial decisions (column 5) varies significantly across KCs: from as few as four on KC_1 to more than 80 on KC_{20} . The average number of tutorial decisions on elicit/tell (ET) (column 3) and justify/skip-justifys (JS) (column 4) also varies across KCs. There are only 4.05 ET decisions on KC_1 and more than 70 on KC_{20} . Similarly, there are only 3.34-3.97 JS decisions for KC_{22} on average and more than 16 for KC_{21} . Overall, the ET tutorial decisions were much more frequent than the JS ones.

6.3 KC-BASED REWARD

In Study 2, the student's final reward was based upon his/her KC-based NLG. More specifically, for each KC_k , the students were divided into two groups, *low learners* and *high learners*, according to a median split of the students' KC-based NLGs. The high learners were assigned a final reward of +100 while the low learners were assigned a final reward of -100.

However, there were at least two limitations from doing this. First, there was little to

Table 6.1: Compare Three Corpus on Eight Primary KCs

		ET	JS	overall	I-ratio	J-ratio	pre-post
KC_1	Exp	4.05		4.05	1.16		$t(126) = 3.28, p = 0.0013$
	Dich	4.16		4.16	2.37		$t(72) = 2.80, p = 0.0066$
	Comb	4.09		4.09	1.33		$t(200) = 4.30, p = 0.0000$
KC_{14}	Exp	7.95	2.00	9	1.36	0.74	$t(126) = 1.076, p = 0.284$
	Dich	8.38	2	9.00	7.44	0.95	$t(72) = 2.03, p = 0.0462$
	Comb	8.11	2	9.00	3.49	0.80	$t(200) = 2.04, p = 0.0422$
KC_{20}	Exp	72.59	15.36	81.53	1.04	1.21	$t(126) = 5.38, p = 0.0000$
	Dich	72.43	16.43	82.84	0.33	0.78	$t(72) = 4.45, p = 0.0000$
	Comb	72.53	15.75	82.01	0.78	1.05	$t(200) = 6.94, p = 0.0000$
KC_{21}	Exp	33.63	16.92	60.75	1.01	1.45	$t(124) = 3.93, p = 0.0001$
	Dich	49.00	16.86	59.05	1.90	0.99	$t(72) = 4.30, p = 0.0001$
	Comb	39.26	16.9	60.13	1.77	1.28	$t(198) = 5.73, p = 0.0000$
KC_{22}	Exp	30.7	3.34	32.06	1.02	1.34	$t(126) = 2.39, p = 0.0184$
	Dich	31.32	3.97	32.78	0.97	1.73	$t(72) = 2.41, p = 0.0185$
	Comb	30.93	3.57	32.33	1.00	1.50	$t(200) = 3.30, p = 0.0011$
KC_{24}	Exp	52.7	15.48	60.45	1.05	1.52	$t(124) = 3.92, p = 0.0001$
	Dich	48.97	15.89	60.24	0.40	0.58	$t(72) = 3.11, p = 0.0027$
	Comb	51.34	15.63	60.38	0.81	1.18	$t(198) = 4.99, p = 0.0000$
KC_{27}	Exp	21.45	4.83	23.89	1.10	1.32	$t(126) = 4.52, p = 0.0000$
	Dich	21.92	4.62	23.73	0.89	1.75	$t(72) = 1.99, p = 0.051$
	Comb	21.62	4.75	23.83	1.02	1.47	$t(200) = 4.67, p = 0.0000$
KC_{28}	Exp	14.2	4.06	16.58	1.06	1.58	$t(126) = 1.911, p = 0.058$
	Dich	14.19	5.14	18.32	3.58	0.69	$t(72) = 1.82, p = 0.07$
	Comb	14.20	4.46	17.22	1.98	1.24	$t(200) = 2.61, p = 0.0099$
Overall	Exp	273.89	56.61	305.48	0.99	1.19	$t(126) = 3.32, p = 0.0012$
	Dich	270.54	58.43	307.57	0.79	0.77	$t(72) = 2.92, p = 0.0046$
	Comb	272.66	57.28	306.25	0.92	1.03	$t(200) = 4.40, p = 0.0000$

no differentiation between the students who learned much more than the median split and those who were merely above the median and between those who were just below the median and those who achieved a much lower score than the median. Additionally, the difference between the high and low learners rewards were always 200 across all KCs, but the actual NLG difference between the high learners and low learners varied across KCs: for example it was 0.30 on KC_{20} but 0.49 on KC_{23} . It is actually difficult to compare the induced KC-based tutorial tactics across KCs (when there were conflicting policies in multi-KC steps) in this way, because it also depends how much difference existed between the high and low learners on that KC. So in Study 3, instead of using a median split, the final rewards were made directly proportional to the real NLG scores.

As described above, one primary goal in this study was to compare the NormGain tutoring tactics with the InvNormGain ones. For inducing NormGain policies the final reward value was set for each d_i on KC_k as: $NLG_{KC_k} \times 100$. That is, the student's KC-based normalized learning gain for the given KC multiplied by 100. For the KC-general policies, the final reward for each super-dialogue d_i was cumulative KC-based $NLG \times 100$ where NLG was the students' learning gain as calculated based on his/her cumulative KC scores. Because $NLG \in (-\infty, 1]$, the maximum final reward was +100 and the minimum was $-\infty$ for enhancing learning tutorial tactics. Therefore, the NormGain tutoring tactics were expected to enhance students' learning.

For inducing InvNormGain policies, the inverted final rewards were used. More specifically, for KC-specific policies the reward was set for each d_i as $(1 - NLG_{KC_k}) \times 100$. For the KC-general tutorial tactics, the reward was set for each d_i as: $(1 - NLG) \times 100$. Because $(1 - NLG) \in [0, +\infty)$, the maximum final reward was $+\infty$ and the minimum was 0. So the induced InvNormGain tutoring tactics were expected to enhance the reward for tutorial actions that contributed less or nothing to the students' learning.

6.4 STATE REPRESENTATION

As described in Chapter 3, the issue of state representation can be divided into four sub-issues. They are 1) defining the potential feature choices in state representation; 2) capping the number of features included in each policy; 3) discretizing the features appropriately; and 4) determining feature selection procedures. Compared to Study 2, several changes were made in state representation. The number of features was expanded, and the maximum number of features that could be included in a policy were also increased. In addition, a different method on feature discretization was adopted, and more general feature selection approaches were explored. How these four sub-issues were addressed in Study 3 is discussed below.

6.4.1 Sub-issues 1: Feature Choices

In Study 2, 18 features were defined in four categories. One of major concern in Study 2 was that the 18 feature choices might not represent the state well enough. For example, all three autonomy features were based on the number of elicits or tells the tutor gave to that point in the session. However, the number may also depend on how much a student said so far rather than the number of times the student input. Two tutorial dialogues can have the same number of elicit/tells, but a student who generated a lot of words in his/her entry generally did more work than another student who only generated one or two words per turn. Therefore, the first motivation in Study 3 was to include more features in each category so that it would represent the dialogue states better.

The second motivation was to expand the number of categories. In addition to the four categories proposed in [Moore et al., 2004], two other categories were included that had been suggested by the previous literature. For example, previous research indicated that there was the learning difference between genders [Coley, 2001, Gallagher, 2001, Quek et al., 2002]. Additionally, we have shown that other background information such as MathSat score can predict a student's learning in math and science [Chi and VanLehn, 2008]. Therefore, a new category of features that included certain background information was added. The category

was named Background Features.

Additionally, previous analyses by Litman’s group have shown that simple linguistic features computed from the students’ contributions to the tutorial dialogue are correlated with learning. Forbes-Riley et al. [Forbes-Riley et al., 2007], for example, discovered that the number of times a student mentioned a physics concept and the number of physics concepts involved in a student’s dialogue were significantly correlated with learning. Additionally, in [Purandare and Litman, 2008] the authors identified several additional features that can be used to predict learning gains. These include the number of physics concepts mentioned in the student’s turn, the concept-to-word ratio, the number of student turns with physics concepts, and so on. Therefore, a new feature category was added that describes the characteristics of dialogue generated by students. This category was named Student Dialogue Features.

In a word, feature choices were expanded from four categories and eighteen features in Study 2 to six categories and fifty features in Study 3. The categories are: amount of the work that the tutor has let the student perform (*Autonomy*); time-related tutorial contextual information (*Temporal situation*); contextual information about the solution process (*Problem Solving Contextual features*); the student’s current performance (*Performance*); background information about the student (*Background*); and semantic information about the students’ tutorial dialogues (*Student Dialogue*). All of these features are static information, or can be computed in real time as the student works. The individual features are described below.

6.4.1.1 Autonomy — five features Autonomy Features relate to the amount of work performed by the student in the dialogue. All five autonomy features end with an ‘A’ in their name and are numeric. Three of the five were included in Study 2 while two of the features, `stuWordsToTuWordsA` and `stuWordsToTuWordsSessionA`, are new. In the following, the label “**” is used to represent that the feature is a new feature and was not included in the state choice in Study 2.

1. **tellsSinceElicitA**: The number of tells the student has received since the last elicit prompt, irrespective of the KC involved. This feature reflects how active a student is right now.

2. **pctElicitA**: The percentage of elicit/tell decision points compared to what the tutor has opted to elicit during the dialogue, irrespective of KC. This feature reflects how active a student is overall.
3. **stuWordsToTuWordsA**** : The ratio of student-generated words to tutor-generated words over the entire tutoring history, regardless of KCs. This feature also reflects how active a student is overall, but it uses the words ratio. This is because when two students receive the same percentage of elicits, a student with higher stuWordsToTuWordsA is assumed to be more active than the one with lower stuWordsToTuWordsA.
4. **stuWordsToTuWordsSessionA****: The ratio of student-generated words to tutor-generated words in this session regardless of KCs. This feature also reflects how active a student is in this session by using the words ratio between the student and the tutor.
5. **pctTellsKCSessionA**: The percentage of tells received this session for the given KC, KC_k . This feature reflects how active a student is on a specific KC in this session.

6.4.1.2 Temporal Situation — three features Temporal Situation Features encode time-related information about the problem-solving process. All three temporal situation features end with a ‘T’ and are numeric. All three were included in Study 2.

1. **durationKCBetweenDecisionT**: Time since the last tutorial decision was made on the current KC. This feature reflects how active a student’s knowledge of the current KC is. If “durationKCBetweenDecisionT” is high, it means that the tutor has not mentioned the KC recently so the student’s knowledge on the current KC may be still.
2. **TimeInSessionT**: The total time spent in the current session. This feature reflects a student’s fatigue level.
3. **TimeBetweenSessionT**: The time elapsed between the end of the previous session and the beginning of the current one. This feature reflects how likely a student has forgotten what they learned in previous sessions.

6.4.1.3 Problem Solving Contextual — fifteen features Problem Solving Contextual features encode information about the current problem-solving context. All fifteen problem solving-related features end with ‘PS.’ In Study 2, we included five features in this

category (The first five features listed below — without “** ” in their names). However, there are certain features that are important to describe the context of the tutorial decisions which were not previously included. Thus, this category primarily consists of 10 new feature choices that the author believes will represent the tutorial context. The selection of these features was informed by prior research. For example, previous research suggests that whether to tell or to elicit should depend on the student’s current competence and how difficult the knowledge is. Therefore, a new feature was included, “conceptDifficultyPS”, to describe the tutorial questions’ difficulty level. Moreover, [Purandare and Litman, 2008] found that the number of concepts introduced by the tutor per-turn correlates with students’ learning, so the features tutAverageConceptsPS** and tutAverageConceptsSessionPS** were in the following list:

1. **EarlyTrainingPS:** For the first three problems, the value is 0 and for the later four problems, the value is 1. This feature reflects how well a student might get used to the tutoring system.
2. **SimpleProblemPS:** The first three problems are categorized as simple problems since solving them involves only one domain principle; the next two are medium; and the final two are complex. This feature reflects the complexity of the training problems’ problem solutions.
3. **DuringWalkThroughPS:** For each training problem, the tutorial dialogues followed a two-stage procedure: first problem solving followed by a post-problem discussion. This feature describes whether a tutorial decision occurred during the problem solving or post-problem discussion.
4. **nKCsPS:** The number of times the present KC has occurred in the current tutorial dialogue. This feature reflects overall how familiar the student is with the current KC.
5. **nKCsSessionPS:** The occurrences of the current KC in the tutorial dialogue in this session so far. This feature reflects how many times the student has accessed the current KC in this session.
6. **newLevelDifficultyPS**:** If the current problem is more complex than the prior problem (i.e. we have crossed a boundary from simple to medium or medium to complex).

This value is 1 for P1, P4, and P6 and 0 for the rest. If a problem is much more difficult than its predecessors, it might take a student long time to learn the problem.

7. **conceptDifficultyPS****: The current question’s difficulty level. This feature is roughly estimated from the combined training corpus of Exploratory and DichGain Corpus. For each tutorial decision step, we count the number of times the tutor decided to elicit the answer from the students from the combined corpus and represent it as *#elicit*. Then among all these occurrences, we count the number of occurrences of correct answers in the corpus and represent it as *#correct*; the number of occurrences of incorrect answers as *#incorrect*; the number of occurrences that students’ simply did not answer the tutor’s question by input “I do not know” as *#unknown*, and finally, the number of occurrences of partially correct answers as *#partial*. A partial correct answer refers to a correct but incomplete answer; for example, to calculate the value of the Kinetic Energy of the rock at T_0 in P4, sometimes a student’s answer was “1.2” instead of “1.2J”. The conceptDifficultyPS was calculated by

$$\text{conceptDifficulty} = \frac{0.0 * \#correct + 1.0 * (\#incorrect + \#unknown) + 0.5 * \#partial}{\#elicit} \quad (6.1)$$

ConceptDifficulty is always in the range of [0, 1]. If conceptDifficulty =1, it means it is a difficult question, whereas if it is close to 0, it means it is an easy question.

8. **QuantitativeDegreePS****: This feature measures how quantitative the tutorial action is. When the value is 1, it indicates the tutorial action is purely quantitative; when it is 0, it is purely qualitative. When it is in the middle, then it is mixed. For example, Line 4 in Figure 1.2 is a quantitative step since the tutor asks about the name of the principle to apply to solve for KE_0 , while Line 6 in Figure 1.2 is a qualitative step since the tutor asks the student “Would the direction of v_0 affect the rock’s kinetic energy at T_0 ?”. For some decision steps, both types of discussion are involved. For example, in a post-problem discussion, the tutorial decision step was about what physics quantities the Kinetic Energy depends on. A correct answer should look like From $KE = \frac{1}{2}mv^2$. We can infer that the Kinetic Energy of an object at time point T is influenced by the mass of the object and its magnitude of velocity at T. This is both a quantitative step in that

students need to know the formula of kinetic energy and it is also a qualitative step in that they need to know what the variables represent here, especially that v only refers to the object's magnitude of the velocity at time T , which does not include the direction. So for this tutorial decision step, the QuantativeDegreePS is 0.5.

9. **numPhysConceptsTutorDialogueSessionPS****: The number of tutor's physics concepts ¹ in this session so far. These physics concepts were identified and generated by the two domain experts and knowledge representations (not the author); these are the key words and physics concepts in the domain of work and energy.
10. **tutAverageConceptsPS****: The average number of a tutor's physics concepts in each turn. This feature reflects how many physics concepts the tutor has mentioned so far and how important the tutor's turns might be. The more frequently physics concepts showed in tutor's turns, the more likely students might learn from these previous tutor's turns.
11. **tutAverageConceptsSessionPS****: The average number of physics concepts in each tutor's turn in this session. This feature reflects how important the tutor's turns might be in this session.
12. **tutConceptsToWordsPS****: The ratio of physics concepts to the words that have been used in the tutor's turn. This feature also reflects how often the tutor has mentioned physics concepts overall.
13. **tutConceptsToWordsSessionPS****: The ratio of physics concepts to the words that have been mentioned in the tutor's turn in this session. This feature also reflects how often the tutor has mentioned physics concepts in this session.
14. **tutAverageWordsPS****: The average number of words in the tutor's turn. This feature reflects how verbose the tutor is overall.
15. **tutAverageWordsSessionPS****: The average number of words in the tutor's turn in this session. This feature reflects how verbose the tutor is in the current session.

¹“Physics concepts:” A word is a physics concept if it is one of the following: 'scalar', 'vector', 'mass', 'displacement', 'velocity', 'acceleration', 'gravitation', 'gravity', 'force', 'weight', 'normal', 'friction', 'system', 'isolated', 'non-isolated', 'kinetic', 'energy', 'potential', 'total mechanical', 'gravitational', 'spring', 'tme', 'spe', 'ke', 'gpe', 'conservation', 'non-conservation', 'work', 'network', 'net', 'direction', 'perpendicular'.

6.4.1.4 Performance — twelve features Performance Features describe information about the student’s performance during the training. All twelve performance-related features end with “PM.” In Study 2, seven features were included (listed as the first seven in below). Five of the twelve features in this category are new for Study 3. Most of the original seven features defined here described a student’s overall performance to that point. However, a more accurate description of his/her performance should be based upon a student’s more recent capability or performance. Therefore, five performance features have been added which measure students’ more recent performance — in the session so far.

1. **pctCorrectPM:** We compute this by assessing all of the correct KCs in students’ entries divided by the total number of KCs in the students’ entries. This feature reflects the student’s overall competence when only elicits are counted as learning opportunities.
2. **pctOverallCorrectPM:** We compute this by assessing all of the correct KCs in student’s entries divided by the total number of KCs shown in both the tutor’s entries and the student’s entries. This feature reflects the student’s overall competence, when both elicits and tells are counted as learning opportunities.
3. **nCorrectKCPM:** The absolute number of correct responses on the current KC in the student’s entries. This feature reflects the student’s overall competence *on the current KC* by measuring how many times the student have given correct responses on the KC.
4. **pctCorrectKCPM:** We compute this by assessing all of the correct cases on the present KC in the student’s entries divided by the total number of cases the present KC showed in the student’s entries. This feature reflects the student’s overall competence on the current KC when only elicits on the KC are counted as learning opportunities.
5. **pctOverallCorrectKCPM:** We compute this by assessing all of the correct cases on the present KC in the student’s entries divided by the total number of cases on the present KC in both the tutor’s and the student’s entries. This feature reflects the student’s overall competence on the current KC when both elicits and tells that involve the KC are counted as learning opportunities.
6. **pctCorrectKCSessionPM:** We compute this by assessing all of the correct cases on the present KC in the student’s entries in this session divided by the total number of cases on the present KC in the student’s entries in this session. This feature reflects the

student's competence on the current KC in this session when only elicits on the KC in this session are counted as learning opportunities.

7. **nIncorrectKCPM**: The number of incorrect student responses on the current KC since the start point. This feature reflects the student's overall incompetence on the current KC.
8. **nCorrectKCSessionPM****: The absolute number of correct responses on the current KC in the student's entries in this session. This feature reflects the student's incompetence on the current KC in this session.
9. **pctCorrectSessionPM****: We compute this by assessing all of the correct KCs in the student's entries in this session divided by the total number of KCs in the student's entries in this session. This feature reflects the student's overall competence across all KCs in this session when only elicits in this session are counted as learning opportunities.
10. **pctOverallCorrectSessionPM****: We compute this by assessing all of the correct KCs in the student's entries in this session divided by the total number of KCs shown in both the tutor's entries and the student's entries in this session. This feature reflects the student's overall competence across all KCs in this session, when both elicits and tells in this session are counted as learning opportunities.
11. **pctOverallCorrectKCSessionPM****: We compute this by assessing all of the correct cases of the present KC in the student's entries in this session divided by the total number of cases of the present KC in both the tutor's and the student's entries in this session. This feature reflects the student's overall competence on the current KC in this session when both elicits and tells that involve the current KC in this session, are counted as learning opportunities.
12. **nIncorrectKCSessionPM****: The number of incorrect student responses on the current KC in this session. This feature reflects the student's overall incompetence on the current KC in this session.

6.4.1.5 Background — five features As described above, previous research has shown that certain background features describe general information about the student's ability to learn. Five Background Features, such as the student's pre-test scores, have been included.

None of these features change during problem solving. All five background features end with “BG.” All features in this category are new features which were not incorporated into Study 2. One important note was that for DichGain group, the following features, genderBG, ageBG, MathSatBG, and VerbalSatBG, were not available because of the administrative error.

1. **genderBG****: The student’s gender. It may be the case that the effectiveness of tutorial policies depends upon differences in gender. For example, male students might learn better by answering the questions, while female ones might learn better by reading the information.
2. **ageBG****: The student’s age. This feature reflects how much school experience the participant might have.
3. **MathSatBG****: The student’s math SAT scores. This feature reflects the participant’s math skill since the physics domain is a math-related domain.
4. **VerbalSatBG****: The student’s verbal SAT scores. This feature reflects the participant’s reading skill since the selected domain also has a lot qualitative discussions.
5. **pretestBG****: The student’s pre-test scores. This feature reflects the participant’s competence in physics before he/she starts the training session.

6.4.1.6 Student Dialogue — ten features This is also a new category. It describes the characteristics of the entries input by students. These are simple linguistic features that are computed from the student’s part of the tutorial dialogues. These features were inspired by previous work on tutoring. Forbes-Riley et al., for example, discovered that the number of times a student mentioned a physics concept and the number of physics concepts involved in the students’ dialogue were significantly correlated with learning [Forbes-Riley et al., 2007]. Additionally, Purandare and Litman, identified several additional features that can be used to predict learning gains: the number of physics concepts mentioned in the students’ turn, the concept-to-word ratio, the number of the student’s turns with physics concepts, and so on [Purandare and Litman, 2008].

1. **averagePhysConceptsStudentDialogueSD****: The average number of physics concepts mentioned per student turn since the training started. This feature reflects how physics-like the student-generated answers have been since the beginning of the training.
2. **numStudentConceptualDialogueSD****: The number of the student's turns that includes at least one physics concept. This feature reflects how many times the student-generated answers included at least one physics concept.
3. **stuConceptToWordRatioSD****: The ratio of physics concept words to total words in the student's turns. This feature also reflects how physics-like the student-generated answers have been since the beginning of the training.
4. **stuAverageWordsSD****: The average number of words per student turn. This feature also reflects how verbose the student was overall. It might also reflect how active the student was.
5. **stuAverageConceptSD****: The average number of the student turns that involve at least one physics concept. This feature reflects how often the student's answers involved at least one physics concepts since the start of the training.
6. **averagePhysConceptsStudentDialogueSessionSD****: The average number of physics concepts mentioned per student turn in this session. This feature reflects how physics-like the student-generated answers are in this session.
7. **numStudentConceptualDialogueSessionSD****: The number of the student turns that mention physics concepts in this session. This feature reflects how many times the student-generated answers included at least one physics concept in this session.
8. **stuConceptToWordRatioSessionSD****: The physics concepts to words ratio per student's turn in this session. This feature also reflects how physics-like the student-generated answers are in this session.
9. **stuAverageWordsSessionSD****: The average length of student turns in this session. This feature reflects how verbose the student was in this session, and it might also reflect how active the student was in this session so far.
10. **stuAverageConceptSessionSD****: The average number of student turns which involve at least one physics concept over all the student turns in this session. This feature

reflects how often the student’s answers involved at least one physics concepts in this session.

6.4.1.7 Simplified Example of Deriving Fifty Features from Log Files. Similar to Chapter 5, the same sample tutorial dialogue in Table 5.1 were transformed into fifty features for inducing tutorial tactics on KC_{20} . I attach how the 50 features were updated as the sample dialogue in Table 5.1 goes on in the Appendix L.

6.4.2 Sub-issues 2: Maximum Number of Features

In Study 2, the maximum number of features was capped at four because of the four categories. The effect of this was that the maximum number of features involved in the induced tutorial tactics was limited to only one. It is quite possible, however, that for some KCs there was more than one feature from one category that should have been included in the state representation. Therefore, in Study 3, no limit was set for the number of features that each category could contain, nor was a requirement set that there needed to be a feature from each category.

In order to determine the maximum number of features in the induced policy, it is necessary to consider the amount of available data and available computational power. In the worst case scenario, there were only 2 JS tutorial decision steps in the DichGain training corpus for KC_{14} . In order to learn effective tutoring tactics, we should have a corpus that covers each of these states at least once. Therefore, based on the minimum data available from the three training corpus for KC_{14} , we capped the number of features in each policy at six, which means that there are at least $2^6 = 64$ states in the learned policy. Alternatively, we could have used a flexible number for different KCs. However, given that six features would already result in very subtle policies, as shown in Table 6.5, it is not the case that learned tutorial tactics with six features were most effective. Instead the final induced policies primarily have 3-5 features in their state representation, and only two of 34 final tactics have six features. So it appears that six is a reasonable number for this study.

6.4.3 Sub-issues 3: Feature Discretization

Five of the fifty features, EarlyTrainingPS, SimpleProblemPS, DuringWalkThroughPS, genderBG, and newLevelDifficultyPS are discrete. The remaining forty-five features are numeric and must be discretized before a suitable MDP can be constructed. Previously, in Study 2 a median split was implemented. For the present study, a more complicated procedure was adopted. The discretization procedure in Study 3 used two clustering procedures, one based upon bounding the number of clusters, and the other based upon identifying the optimal cluster means.

For each of the continuous valued features, the ideal number of clusters was identified using a TwoStep package embedded in SPSS. TwoStep clustering is a scalable cluster analysis algorithm designed to handle very large datasets. It is capable of handling both continuous and categorical variables and attributes. Its key advantage is that it can find the optimal number of clusters when the ideal value is unknown. Once the proper number of clusters has been determined, it is possible to apply more traditional K-mean methods to identify the contents of each cluster. After the clusters and their mean values were identified, the clusters were ranked by value, assigned the student values, and discretized according to cluster membership.

Although the median split in study 2 may not have been optimal, it capped the number of possible values for each state feature at two. By using the automatic procedure, it is highly likely that each feature it would have been discretized into many clusters. Increasing the number of possible values in a state would increase the number exponentially. For example, the application of the TwoStep package to stuAverageConceptSD from the Exploratory training corpus on KC_{20} resulted in seven clusters. If each feature has seven possible values, for a six-feature policy, it would have 7^6 . In order to control the number of possible values, the maximum number of discrete values for each feature was set at four. However, this procedure is still risky in that it can generate small clusters containing at most a handful of data points. This can then lead to problems of data sparsity during the RL phase. These two problems were addressed by adding a reduction loop to the procedure. If any given cluster within a set contained less than 20% of the total cases after K-means clustering took

place, then the number of clusters was reduced by one and the K-means algorithm was run again. This process was repeated until all clusters exceeded the 20% threshold.

If this requirement could not be met, then the feature was discretized through a median split. A Pseudo-code representation of the algorithm is shown below:

1. for each feature choice f_i :
2. **Step 1:** Count the clusters $numCluster$ by applying two-step clustering approach
3. **Step 2:** Via k-means clustering, discretize the feature f_i into $numCluster$ clusters.
4. **Step 3:** if one cluster has less than 20% cases:
5. **Step 4:** if $numCluster = 2$:
6. **Step 5:** median split
7. **Step 6:** else:
8. **Step 7:** $numCluster = numCluster - 1$
9. **Step 8:** Go to Step 2.

For example, by running this procedure, the system performed a median split on “stu-AverageConceptSD” so that values in the range of 0 to 0.228395 is 0 and values in the range of 0.228395 to 1 is 1 (The feature choices were normalized in this dissertation). Finally, the number of clusters for each feature ranged from, at minimum, two clusters to as many as four.

6.4.4 Sub-issues 4: Feature Selection

In Study 2 a simple greedy-feature selection method was used. The procedure was repeated later, using the same Exploratory Corpus, the same 18 features and the discretization procedure, rewards and so on. The only difference was that some new feature selection methods were applied. Results showed that by simply changing the feature selection methods, the induced policies had a much higher ECR than the ones used in Study 2. Recall that the ECR was the criteria for picking the best policies [Chi et al., 2008a]. For example, the policy $\pi(KC_{22}, ET)$ used in study 2 previously had an ECR of 9.40. Under the new feature selection methods, the policy $\pi(KC_{22}, ET)^*$ had an ECR of 44.29, almost four times higher.

Therefore, in Study 3 a more complex set of eleven feature-selection approaches were explored. The domain general feature selection methods were the main ones explored, and by doing so, the relationships among the features were neglected. For example, a domain-specific feature selection approach could only select features that correlated with NLG. However, by

doing so, most of the problem-solving contextual features would never be considered such as DuringWalkThroughPS, EarlyTrainingPS, and so on.

Almost every approach described below involved inducing single-feature policies first. That is, for each of the fifty feature choices, the RL package was used to induce a single-feature-policy. Because generating single-feature policies does not involve any feature selection procedure, such policies were labelled as “single”. In the following, the focus is on using feature selection to select at least two features in a policy. To differentiate from the single-feature policies, the policies induced through feature selection were labelled as non-single-feature policies. In short, Study 3 explored eleven feature-selection methods to induce non-single-feature policies. Four of the approaches were based upon RL (*Upper_Bound*, *Lower_Bound*, *ECR*, and *Hedge*) used in the previous studies [Chi et al., 2008a]; one was based on Principal Component Analysis (PCA) (*PCA-only*); four were combinations of PCA and RL (*PCA-Upper_Bound*, *PCA-Lower_Bound*, *PCA-ECR*, and *PCA-Hedge*); while the final pair were based upon stochastic selection (*Random* and *PCA-random*).

In the following, $\pi(KC_i, D_j, NormGain)$ and $\pi(KC_i, D_j, InvNormGain)$, are used to represent a NormGain and an InvNormGain KC-based policy on KC_i for tutorial decisions D_j respectively. Here $D_j \in \{ET, JS\}$, KC^* is used to represent KC-general policies, and $KC_i \in \{KC_1, KC_{14}, KC_{20}, KC_{21}, KC_{22}, KC_{24}, KC_{27}, KC_{28}, KC^*\}$.

6.4.4.1 RL-based Feature Selection As described in Chapter 3, after being given a complete MDP structure, the Tetreault and Litman’s toolkit would calculate a policy together with the policy’s ECR and 95% CI [Tetreault and Litman, 2008]. Lower-Bound and Upper-Bound were used to refer to the 95% confidence bounds calculated for the ECR. For example, a final tutorial tactic in Study 3 $\pi(KC_{14}, ET, NormGain)$ was based on feature: durationBetweenDecisionT alone which is derived from the combined corpus. $\pi(KC_{14}, ET, NormGain)$ states that “if the duration since last the tutorial decisions on KC_{14} is less than 160.07 sec, then the tutor should elicit.”; $\pi(KC_{14}, ET, NormGain)$ has ECR = 9.99 (range $[-\infty, 100]$) with a 95% confidence interval = [9.85, 10.06], which means there is a 95% chance that the ECR of the learned policy is between a lower-bound of 9.85 and an upper-bound of 10.06. Another $\pi(KC_{14}, ET, NormGain)^*$ was based on the

feature conceptDifficultyPS and is derived from training the Exploratory Corpus; it states that: “if the current tutorial decision step is easy ($< 49.53\%$), then the tutor should elicit.” $\pi(KC_{14}, ET, NormGain)^*$ has ECR = 1.19 with a 95% confidence interval = [0.03, 3.39].

To this point ECR has always been used as the criteria for selecting the best policies. However, the policy’s Lower-Bound or Upper-Bound can also be used as the criteria. More specifically, the former evaluates the performance of policies in the worst case, while the latter describes how well the policy can perform. As in the example above, $\pi(KC_{14}, ET, NormGain)$ is more effective than policy $\pi(KC_{14}, ET, NormGain)^*$ because even its Lower-Bound is much higher than policy $\pi(KC_{14}, ET, NormGain)^*$ ’s Upper-Bound. Sometimes researchers encounter situations in which the ECR for Policy A is the same as the ECR for Policy, B, but the confidence interval of A is much narrower than that of B. In this case, a new criterion, Hedge, can be applied to compare the two policies. Hedge is defined as a learned policy:

$$Hedge = \frac{ECR}{UpperBound - LowerBound} \quad (6.2)$$

By applying Hedge, policy A is shown to be more effective than policy B. Any of these criteria, ECR, Lower-Bound, Upper-Bound, or Hedge can be used to evaluate policies, and thus they are used as four different criteria for feature selection. These feature-selection methods are fairly straightforward and use the same general procedure, described below.

For each ranking metric in [ECR, Lower-Bound, Upper-Bound, Hedge]

1. For each of the 50 feature choices, use the RL package to induce a single-feature-policy.
2. Rank the policies in descending order based upon the ranking metric.
3. For $i = 2$ to 6
 - Pick the top i features from the ranked list and construct an MDP using them for the state representation.
 - Induce a policy for that MDP and set it aside.

These feature-selection metrics based upon the sorting criteria used: ECR, Lower-Bound, Upper-Bound, and Hedge respectively. These are RL-based methods as the feature selection procedures all use the policy ECR to identify optimal feature choices. For each of the ranking metrics, the above procedure resulted in five lists of policies and resulted in twenty policies

that involved at least two features in the state representation for each KC on each type of tutorial action decision from each corpus.

Here is an example which applies the Upper-Bound feature selection for deriving policies on KC_{14} for the ET decision from the Exploratory Corpus. First, from the fifty features fifty single-feature policies were learned: π_1, \dots, π_{50} . Then the fifty single features were sorted based on the upper-bound of its corresponding single-feature policy. Then the first six features were selected for which the corresponding single-feature policies had the highest upper-bounds. In this example, the six selected features are in the order of frequency: {durationBetweenDecisionT, numPhysConceptsTutorDialogueSessionPS, nKCsPS, nKCsSessionPS, nIncorrectKCSessionPM, tutAverageWordsSessionPS}. For the first feature, “durationBetweenDecisionT” a single-feature policy has already been learned, so the process begins with the second one. The process used is to learn a two-feature policy based on the first two features: “durationBetweenDecisionT, numPhysConceptsTutorDialogueSessionPS”, and then a three-feature policy based on the first three features: “durationBetweenDecisionT, numPhysConceptsTutorDialogueSessionPS, nKCsPS”, and so on. The most complicated policy to be learned would be a six-feature policy which includes all six features in the state. Therefore, for KC_{14} on ET tutorial decisions, five non-single-feature tutorial tactics were induced by following the Upper-Bound feature from each training corpus.

6.4.4.2 PCA-based Feature Selection Given the data sparsity problems that most machine learning techniques face, an ideal state representation should have as few features as possible while still being rich enough to represent the domain accurately. Unfortunately, some of the features available in this study were highly correlated which reduced their expressiveness when used together. For example, given $nCorrectKCPM$, the number of correct responses on the current KC KC_k , and $nIncorrectKCPM$, the number incorrect responses on KC_k generated by the student, $pctCorrectKCPM$, the percentage of correct responses given by the student on KC_k could easily be calculated as: $pctCorrectKCPM = \frac{nCorrectKCPM}{nCorrectKCPM+nIncorrectKCPM}$. Therefore, it was necessary to apply an analysis procedure to avoid redundant features. One such procedure explored in this thesis is Principal component analysis (PCA) [Jolliffe, 2002].

PCA is a mathematical procedure that transforms a large number of, possibly correlated, variables into a smaller number of uncorrelated variables called principal components. It is a popular dimensionality-reduction technique as it is simple, non-parametric, unsupervised, and has been applied successfully in a number of domains. More formally, given a list of $n - dimension$ variables, PCA extracts a reduced set of p ($p < n$) principal components or factors that account for most of the variance present in the original set.

This is done by first extracting n principal components from the original n variables. Each component is a linear combination of the variables. Formally, it results in a set of n linear equations with n unknown variables. In other words, PCA repackages the original variables into an equal number of uncorrelated principal components. The first of these components accounts for the largest possible amount of variance. The second component, which attempts to explain the variance remaining after the first component has been extracted, accounts for the second largest amount of variance, and so on. As the components are extracted they are restricted to be orthogonal. Geometrically, they may be viewed as defining an n -dimensional space.

The variance in this correlation matrix is “repackaged” into a set of n eigenvalues, one for each principal component. Each eigenvalue represents the amount of variance covered by its associated component. Thus the first eigenvalue accounts for the largest share of the variability with each one accounting recursively for the largest share of the *remaining* variability once its predecessors are accounted for. Each of the n variables’ variance is normalized to 1. Each component’s eigenvalue may be compared to this standard value to determine how much more or less variance it accounts for than a single variable. With n variables there is a total of n variance to distribute. The extracted components, like the set of variables, account for all of this variance collectively. The proportion of variance accounted for by one component equals its eigenvalue divided by n .

Typically, the goal was to derive a set of less than n components. When a set of p out of the variables share a considerable amount of variance, then $p < n$ components will have relatively large eigenvalues, while the rest will have substantially smaller eigenvalues. In this situation what needs to be determined is how many components will be retained and how many will be discarded. One rule of thumb is to drop any component with an eigenvalue

of less than 1, that is, any component that accounts for less variance than a single variable. Having done that, the $n - dimensional$ space defined by the original variables has been reduced to a $p - dimensional$ component space that still covers the bulk of the variance.

In this study, initially, all fifty feature choices were normalized. PCA was then applied to the normalized features to generate fifty principal components and their corresponding eigenvalues. These eigenvalues were arranged in descending order, and all components whose eigenvalues were less than 1 were removed. For each eigenvalue, the feature that was maximally correlated with the corresponding principal component was identified. The resulting features were a subset of the original fifty feature choices that were designated the *PCA-feature subset*. *PCA-feature subset* is an ordered list arranged by the eigenvalues of its corresponding principal components. Results showed that the number of PCA-features selected for each KC varied from eight to thirteen.

Once the *PCA-feature subset* was identified, the PCA-only feature selection procedure was straightforward. It began with the first feature in *PCA-feature subset* and added one feature at a time to the MDP and learned a new policy. This process was repeated five times.

Here is an example. After running PCA on a list of 50 – *dimension* feature variables exacted from the Exploratory Corpus for inducing KC-general tutorial tactics on elicit/tell, a reduced set of thirteen principal components whose eigenvalues were bigger than 1 was extracted. They were ordered by their corresponding eigenvalues as shown in Table 6.2. In that table, the second column shows the eigenvalues of the corresponding components arranged in decreasing order. Specifically, eigenvalues decreased from 6.59 for the first component to 1.11 for the 13th component. The third column represents the cumulative eigenvalues. For example, the last row of the third column is 81.42%, which means these 13 principal components have 81.42% of the total information provided by the original fifty variables.

For each principal component, one feature with the highest correlation with the component was extracted. In this instance, “pretestBG” (the students’ pre-test score) is selected for the first principal component, SimpleProblemPS (whether the current problem students are working on is a simple problem or not) as the second principal component, and so on. The corresponding correlation value is listed in the last column. So our *PCA-*

Table 6.2: An Example PCA Feature Set from the Exploratory Corpus Induced for A KC-general Tutorial Tactics

Order	EigenVals	EigenCumulative	EigenVar	EigenCorr
1	6.589	13.179	pretestBG	0.962
2	4.858	22.894	SimpleProblemPS	0.933
3	4.746	32.387	nKCsSessionPS	0.939
4	4.416	41.218	stuAverageWordsSD	0.912
5	4.304	49.827	tutAverageConceptsSessionPS	0.859
6	3.626	57.08	pctCorrectSessionPM	0.882
7	3.601	64.283	stuConceptToWordRatioSessionSD	0.866
8	2.555	69.393	MathSatBG	0.914
9	1.357	72.107	newLevelDifficultyPS	0.871
10	1.205	74.516	pctElicitA	0.618
11	1.202	76.919	QuantativeDegreePS	0.868
12	1.141	79.202	conceptDifficultyPS	0.578
13	1.11	81.421	durationBetweenDecisionT	0.763

feature subset is simply the collection of the thirteen features in column 4, sorted by the corresponding eigenvalues, which is pretestBG, SimpleProblemPS, nKCsSessionPS, stuAverageWordsSD, tutAverageConceptsSessionPS, pctCorrectSessionPM, stuConceptToWordRatioSessionSD, MathSatBG, newLevelDifficultyPS, pctElicitA, QuantativeDegreePS, conceptDifficultyPS, durationBetweenDecisionT.

For PCA-only feature selection, the researcher started with a single feature policy, the first Eigen Variable pretestBG here, then a two-feature policy by using the first two Eigen Variables: pretestBG, SimpleProblemPS, and so on. The most complicated policy would be a six-feature policy (because the maximum number of features in a policy was capped at six) by using the first six Eigen Variables: pretestBG, SimpleProblemPS, nKCsSessionPS,

stuAverageWordsSD, tutAverageConceptsSessionPS, pctCorrectSessionPM, which included one background feature “BG”, three problem state contextual features “PS”, one feature about the student’s dialogue “SD”, and one feature about the student’s performance “PM”.

Therefore, for each KC for each type of tutorial action from each individual corpus, five non-single-feature tutorial tactics were induced using the PCA-only feature selection.

6.4.4.3 PCA and RL-based Feature Selection Thus far, four RL-based feature selection methods and a PCA-only feature selection method have been described. By simply combining PCA-only feature selection with the four RL-based feature selection methods, four new feature selection approaches are created. In this method, PCA is used to identify the PCA-feature subset from the original fifty features, creating a set of available features that have eigenvalue greater than one. The four RL-based methods, *PCA-Upper_Bound*, *PCA-Lower_Bound*, *PCA-ECR*, and *PCA-Hedge* are then applied as before. In effect these combined feature selection methods are being used to winnow the set of available features, not once, but multiple times.

Here is a summary of the procedure:

[Stage 1:] Select the PCA-feature Subset:

[Phase 1:] Apply PCA on fifty features.

[Phase 2:] Identify the set of principal components with its eigenvalues greater than or equal to one.

[Phase 3:] For each component, identify the feature F that is most correlated with the component.

The resulting features are the PCA-feature choices and are ranked in the order of their eigenvalues.

[Stage 2:] Apply the RL-based approach:

For each ranking metric in [ECR, Lower-Bounds, Upper-Bounds, Hedge]

For each of PCA-feature choices, use the RL package to induce a single-feature-policy.

Rank the policies in descending order based upon the ranking metric.

For $i = 2$ to 6:

```
Pick the top i features from the ranked list and construct
an MDP using them for the state representation.
Induce a policy for that MDP and set it aside.
```

Based on the sort criteria in phase 2, four feature selection methods were named PCA-ECR, PCA-Lower_Bound, PCA-Upper_Bound, and PCA-Hedge respectively. Similar to previous approaches, for each KC on each type of tutorial action from each individual corpus, five non-single-feature tutorial tactics were induced by combining the PCA and RL-based feature selection methods.

6.4.4.4 Random Feature Selections Thus far nine feature selection methods have been introduced. In order to evaluate their relative effectiveness, a random feature selection method was also introduced. The expectation was that the nine feature selection approaches described above would be more effective than a random feature selection. In other words, it was anticipated that the final tutorial tactics would be induced by the feature selection methods introduced above rather than the random feature selection. Two random feature selection procedures were employed: *Random* and *PCA-Random*. In the former case features were randomly selected from all fifty feature choices. In the latter case PCA-based feature reduction was applied to reduce the set of variables to those with high variance correlations, and then features were randomly selected from the reduced set. Here is the summary of this procedure:

Random-selection :

```
For j= 1 to 2
  For i = 2 to 6
    Randomly select i features from fifty features.
    Induce a policy for that MDP and set it aside.
```

PCA-Random:

[Stage 1:] Select the PCA-feature Subset:

[Phase 1:] Apply PCA on fifty features.

[Phase 2:] Identify the set of principal components with eigenvalues greater than or equal than one.

[Phase 3:] For each component, identify the feature F that is most correlated with the component.

The resulting features are the PCA-feature choices and are ranked in the order of their eigenvalues.

[Stage 2:] Random Selection:

For j= 1 to 2

For i = 2 to 6

Randomly Select i features from the PCA-feature Subset.

Induce a policy for that MDP and set it aside.

For each KC for each type of tutorial action from each individual corpus, 10 non-single-feature tutorial tactics were induced by following either random or PCA-random feature selection.

To summarize, for each KC_i and decision (ET or JS) $\langle KC_i, D_j \rangle$, three training corpora, a space of fifty features, and eleven feature selection methods were explored. For each KC_i and decision pair one set of policy choices was collected for each training corpus. For each corpus, there were fifty single-feature policies. Applying eleven feature selection methods to them yielded $5 \times 9 + 10 \times 2 = 65$ non-single-feature policies (the random and PCA-random feature selection yielded 10 non-single-feature policies each and the remaining nine methods yielded five non-single-feature ones each). A total of 115 potential tutorial tactics were generated for a single KC and decision pair per training corpus. Taken together, all three corpora resulted in a total of $115 \times 3 = 445$ policies for each pair $\langle KC_i, D_j \rangle$. The best policy for $\langle KC_i, D_j \rangle$ was selected from this pool by ECR. For the purposes of this study, the highest ECR irrespective of the confidence bounds or hedging was selected. This is similar to Study 2.

6.5 CONFLICTING POLICIES

In Study 2, a total of 21 KCs for elicit/tell decisions and 10 KCs for justify/skip-justify decisions were considered. When faced with conflicting tutorial policies, the policy with the highest ECR was chosen. However, in the selected domain, certain KCs are necessary precursors for other KCs. KC_{24} , for example, is the definition of Total Mechanical Energy : $TME = KE + GPE + SPE$. In order to apply KC_{24} effectively, students need to know the definition of Kinetic Energy, which is KC_{20} $KE = \frac{1}{2}mv^2$. Thus KC_{20} is a necessary precursor to KC_{24} . In Study 3, we have the ECR of $\pi(KC_{20}, ET) = 14.25$ while the ECR of $\pi(KC_{24}, ET) = 13.51$. In some of the ET tutorial decisions that involve both KC_{24} and KC_{20} , KC_{24} , what would the major topic of discussion be, given this precursor relationship? By following the old procedure on conflicting policies, the system would always choose the policy with the highest ECR. In this case, the system would follow the $\pi(KC_{20}, ET)$ because its ECR is higher than $\pi(KC_{24}, ET)$'s ECR, even though the KC_{24} is the target topic. Therefore, in Study 3 the focus was narrowed eight main KCs and a more complex conflict-resolution approach was adopted. The KCs were first grouped according to three levels of “priority” based upon the domain knowledge with “Top” being the highest and “Low” being the lowest:

Top Level: KC_{27}, KC_{28} .

Medium Level: KC_{24}, KC_{14} .

Low Level: $KC_{20}, KC_{21}, KC_{22}, KC_1$.

This heuristic was built in such way that lower-level KCs are necessary precursors for the higher level KCs. KC_{27} for example is conservation of total mechanical energy : $TME1 = TME0$. In order to apply KC_{27} effectively, students need to know about the definition of Total Mechanical Energy, which is KC_{24} $TME = KE + GPE + SPE$. Thus KC_{24} is a necessary precursor for KC_{27} and KC_{27} is not a necessary precursor for any other KCs in the domain. Therefore we put KC_{27} in the highest level: Top Level and KC_{24} in the second level: Medium Level. By always choosing the highest level KCs in a multiple-KC decision step, we can pick the policy on the target KCs.

When a tutorial decision involved multiple KCs, the system first collected the set of policies at the highest level. If a tutorial decision step did not involve any of the eight major KCs, then the system will follow the KC-general tactics. If, however, a decision step did involve some of the eight primary KCs, then the system would poll the tutorial tactics for KCs at the maximum rank, and will follow the policy with the highest ECR.

6.6 SUMMARY: INDUCTION OF TUTORIAL TACTICS IN STUDY 3

For Study 3 two sets of tutorial tactics were induced. The first is Normalized Gain or the NormGain set. The second is Inverse Normalized Gain or the InvNormGain set. Both sets were induced using the same RL procedures on the same corpora and differed only in the reward function applied to the training corpora. The NormGain set used a positive NLG-based reward and the InvNormGain set used an inverse of that same function.

Both the final NormGain and InvNormGain sets contained 17 policies. Thus, we have a total of 34 policies. In each set, two out of the 17 are KC-general policies, one is an ET policy and the other is a JS policy. The remaining 15 policies are KC-based and consist of seven pairs of policies for seven of the eight main KCs and one policy for KC_1 . Each of the seven pairs is associated with a single KC and contains one ET policy and one JS policy. KC_1 does not arise in any JS decisions and thus only an ET policy was induced for it.

In order to examine a range of possible tactics for each KC_i on either ET or JS tutorial decisions $\langle KC_i, D_j \rangle$, three training corpora, a space of fifty features, and eleven feature selection methods were used to yield 445 policies. As discussed above, the corpora were: the Exploratory Corpus collected in Studies 1, the DichGain Corpus collected in Studies 2, and a combined corpus that merged both datasets. The fifty feature choices could be divided into six categories as described in detail above, and eleven feature selection methods could be applied to them. The best policy for each pair $\langle KC_i, D_j \rangle$ was selected from 445 policies by ECR. For the purposes of this study only the policy with the highest ECR irrespective of the confidence bounds or hedging was selected. This selection process was repeated for each of the 34 policies with the KC-general policies being chosen from models

that ignored the involved KCs. The full list of NormGain policies used in Study 3 are shown in Appendix N and the full list of InvNormGain policies are shown in Appendix O. They detail the policies themselves, the corpus from which they were drawn, the features involved, their discretization ranges, the feature selection method used, the policy’s ECR and its confidence bounds.

The resulting NormGain and InvNormGain policies were implemented back into Cordillera yielding two new versions of the system, named NormGain-Cordillera and InvNormGain-Cordillera respectively. Both systems applied the policies to guide tutorial decisions. KC-based decisions were guided using the KC-based policies. Conflicts were resolved using the ranking among the primary KCs discussed in Section 6.5. For steps that did not involve one of the eight primary KCs, the system used the KC-general tutorial tactics.

Table 6.3 summarizes the major differences between the RL procedure used in Study 2 and that in Study 3. The RL procedures differed on all the major issues. In the next section the discussion will focus on some general characteristics of the induced tutorial tactics. Specifically, it will focus on the source corpus that each of 34 tutorial tactics was derived from, the features that were involved, and which feature selection method yielded the most tutorial tactics, and so on.

Table 6.3: Issue-by-Issue Comparison of Studies 2 and 3

Aspects	Study 2	Study 3	
	DichGain	NormGain	InvNormGain
Training Corpora	Exploratory	Exploratory, DichGain, & Combined	
KCs:	31 KCs	8 main KCs	
Features:	18 features 4 categories	50 features 6 categories	
Discretization:	Medium Split	TwoStep first and then k-means	
Feature Selection	Category-based	11 including Random; ECR; Hedge; & PCA	
Max features /policy:	4	6	
Reward:	$(NLG > median)$ $\rightarrow +100;$ $(NLG \leq median)$ $\rightarrow -100$	$NLG \times 100$	$(100 - NLG) \times 100$
Conflicting Policies:	Follow max ECR	Use KC ranking then ECR.	

6.7 INDUCED POLICIES

In this section, the induced tutorial tactics will be described by identifying the training corpus that each final tutorial tactic was derived from, which feature categories were most frequently involved in the final tutorial tactics, and which feature selection method discovered the most final tutorial tactics. The full list of policies used in Study 3 are shown in Appendix N and Appendix O. The purpose of this section is to determine how RL-related decisions described in the previous section had impacted the induced tutorial tactics. For example, one decision was made to use all three training corpora, did the final induced policies come from one corpus or from all three corpora? Moreover, which features appeared in the final

induced tutorial tactics? Which feature selection method(s) seemed to be more effective? This section begins with a discussion of the training corpus involved in the final 34 tutorial tactics.

6.7.1 Source Training Corpus

Table 6.4 shows which corpus was used to induce the corresponding tutorial tactics. The second and third columns show the source training corpus used in deriving NormGain tutorial tactics on ET and JS for corresponding KCs respectively. The fourth and fifth columns show similar information for the InvNormGain tutorial tactics. The last three rows 10-12 summarize the number of tutorial tactics derived from each corresponding training corpus. For example, Rows 10 and 11 show that the Exploratory Corpus and the DichGain Corpus each generated sixteen final tutorial tactics. The Exploratory Corpus was used to generate 11 NormGain tutorial tactics (5 ET and 6 JS) and 5 InvNormGain ones (3 on ET and 2 on JS) while the DichGain Corpus was used to generate five NormGain tutorial tactics (3 ET and 2 JS) and eleven InvNormGain ones (6 on ET and 5 on JS). The combined corpus, however, only generated one tutorial tactic each for NormGain and InvNormGain.

Table 6.4 also shows that both the Exploratory and DichGain Corpora were involved in generating the final tutorial tactics. However, the majority of the NormGain tutorial tactics were from the Exploratory Corpus, eleven out of seventeen, while most of the InvNormGain tutorial tactics were from the DichGain Corpus, also eleven out of seventeen. This result suggested that the choice of the training corpus is very important for deriving tutorial tactics. However, it is not very intuitive to determine why most of the NormGain policies were from Exploratory Corpus, while most InvNormGain ones were from DichGain Corpus. Future work is needed to explore the characteristics of a training corpus and how to choose a training corpus.

Table 6.4: The Source Training Corpus Of the Inducing 34 Tutorial Tactics

		NormGain		InvNormGain		
		ET	JS	ET	JS	
1	KC_1	DichGain		DichGain		
2	KC_{14}	Combined	Exploratory	DichGain	Combined	
3	KC_{20}	Exploratory	Exploratory	Exploratory	DichGain	
4	KC_{21}	Exploratory	Exploratory	Exploratory	Exploratory	
5	KC_{22}	Exploratory	Exploratory	Exploratory	DichGain	
6	KC_{24}	DichGain	Exploratory	DichGain	DichGain	
7	KC_{27}	Exploratory	Exploratory	DichGain	DichGain	
8	KC_{28}	DichGain	DichGain	DichGain	Exploratory	
9	Overall	Exploratory	DichGain	DichGain	DichGain	Total
10	Exploratory	5	6	3	2	16
11	DichGain	3	2	6	5	16
12	Combined	1	0	0	1	2

6.7.2 Number of Features

Table 6.5 shows the number of the features involved in the thirty-four final tutorial tactics. The second and third columns show the training corpus used in deriving NormGain tutorial tactics on ET and JS for corresponding KCs respectively. The fourth and fifth columns show the same information for the InvNormGain tutorial tactics. Table 6.6 summarizes distribution of sizes, that is, how many policies contained one feature, two, and so on. For example, Row 12 in Table 6.6 shows that there were eight policies that involved three features, four NormGain ones and four InvNormGain ones. To our surprise, only two tutorial tactics involved six features and most of the policies involved three to five features.

Table 6.5: The Complexity of the 34 Induced Tutorial Tactics

#	KC	NormGain		InvNormGain	
		ET	JS	ET	JS
1	KC0	4	5	5	4
2	KC ₁	1		2	
3	KC ₁₄	1	1	1	1
4	KC ₂₀	3	5	3	5
5	KC ₂₁	3	3	6	3
6	KC ₂₂	2	5	3	4
7	KC ₂₄	4	6	2	2
8	KC ₂₇	4	4	4	5
9	KC ₂₈	5	3	3	5

Table 6.6: Distribution of Policy Sizes.

#	Size	NormGain		InvNormGain		Total
		ET	JS	ET	JS	
10	1	2	1	1	1	5
11	2	1	0	2	1	4
12	3	2	2	3	1	8
13	4	3	1	1	2	7
14	5	1	3	1	3	8
15	6	0	1	1	0	2

6.7.3 Feature Choices

The total number of feature occurrences across all thirty-four tutorial tactics was 117. For each induced tutorial tactic, the number of features involved were counted and then totaled. If a feature occurred in several induced tutorial tactics, then each occurrence was counted as one. More specifically, the total number of feature occurrences across the NormGain and InvNormGain tutorial tactics was fifty-nine and fifty-eight respectively.

6.7.3.1 Autonomy Features Autonomy Features relate to the amount of work done by the student in the dialogue. Five Autonomy features were defined: tellsSinceElicitA, pctElicitA, stuWordsToTuWordsA**, stuWordsToTuWordsSessionA**, and pctTellsKCSessionA. As mentioned earlier, features with “**” were new ones added in Study 3. The five autonomy features occurred thirteen times. Among the five features, with the exception of tellsSinceElicitA, the remaining four features occurred only once in one final induced policy. Feature tellsSinceElicit occurred in nine out of the thirty-four final tutorial tactics included in the state representation: five for NormGain and four for InvNormGain. Table 6.7 summarizes the occurrences of each feature in the induced NormGain and InvNormGain tutorial tactics. The number in the parenthesis refers to the number of occurrences. For example, NormGain(8) means there were eight occurrences of autonomy features in NormGain tutorial tactics and pctElicitA (1) means that “pctElicitA” occurred once in the final tutorial tactics. The two new features labelled with “**” occurred only twice.

Table 6.7: Occurrence of Autonomy Features in The Final Tutorial Tactics

		NormGain (8)	InvNormGain(5)
1	tellsSinceElicitA (9)	5	4
2	pctElicitA (1)	1	0
3	stuWordsToTuWordsA** (1)	1	0
4	stuWordsToTuWordsSessionA** (1)	1	0
5	pctTellsKCSessionA (1)	0	1

6.7.3.2 Temporal Situation Features Temporal Situation Features encode the time-related information about the problem-solving process. Three features are defined: `durationKCBetweenDecisionT`, `TimeInSessionT`, and `TimeBetweenSessionT`. Table 6.8 summarize the number of occurrences of each feature in the induced `NormGain` and `InvNormGain` tutorial tactics. The three features occurred a total of fourteen times in the final thirty-four policies. In Table 6.8, Row 1 shows that `durationBetweenDecisionT` showed up eight times, more frequently than the other two features: four times in `NormGain` tutorial policies and four times in `InvNormGain` ones.

Table 6.8: Occurrence of Temporal Situation Features in The Final Tutorial Tactics

		NormGain (7)	InvNormGain(7)
1	<code>durationBetweenDecisionT</code> (8)	4	4
2	<code>TimeBetweenSessionT</code> (2)	1	1
3	<code>TimeInSessionT</code> (4)	2	2

Table 6.9: Occurrence of Problem Solving Contextual Features in The Final Tutorial Tactics

		NormGain (30)	InvNormGain (28)
1	EarlyTrainingPS (2)	1	1
2	SimpleProblemPS (2)	1	1
3	DuringWalkThroughPS (6)	2	4
4	nKCsPS (4)	3	1
5	nKCsSessionPS (3)	2	1
6	newLevelDifficultyPS** (4)	2	2
7	conceptDifficultyPS** (12)	7	5
8	QuantativeDegreePS** (5)	0	5
9	numPhysConceptsTutorDialogueSessionPS**(1)	1	0
10	tutConceptsToWordsPS** (8)	5	3
11	tutConceptsToWordsSessionPS** (4)	3	1
12	tutAverageWordsPS** (5)	2	3
13	tutAverageWordsSessionPS** (2)	1	1

6.7.3.3 Problem Solving Contextual Features Problem Solving Contextual features encode information about the current problem-solving context. There are fifteen features defined in this category. Table 6.9 summarizes the occurrences of each feature in the induced NormGain and InvNormGain tutorial tactics. This category seems to be the most active. There were fifty-eight occurrences in the final thirty-four tutorial tactics, which represents approximately half of all of feature occurrences. In some of the tutorial tactics, more than one feature from this category was involved.

Among the fifteen Problem Solving Contextual features, *conceptDifficultyPS*** (Row 7) is the most frequently occurrences, occurring in twelve induced tutorial tactics: seven for NormGain and five for InvNormGain. The next most frequently occurring feature is in Row

ten: `tutConceptsToWordsPS**`, which describes the ratio of the tutor’s physics concepts to their words. Two features (`tutAverageConceptsPS**` and `tutAverageConceptsSessionPS**`, which represent the average number of tutor’s physics concepts in each turn overall and in this session specifically, did not appear in any of the final induced tutorial tactics. Among the fifty-eight occurrences, new features added for Study 3 occurred forty-one times: twenty-one for NormGain and twenty on InvNormGain ones.

6.7.3.4 Performance Features Performance Features describe information about the student’s performance during problem solving. Twelve feature choices were defined in this category. Table 6.10 summarizes the occurrences of each feature in the induced NormGain and InvNormGain tutorial tactics. The features in this category occurred seventeen times. Row eight shows that “`nIncorrectKCPM`” (the number of incorrect response in the student’s dialogue so far) is the most frequently occurring feature in that it appeared in five final tutorial tactics: two for NormGain and three for InvNormGain. A feature such as `pctOverallCorrectPM` did not appear in any of the final tutorial tactics, probably because the closely-related feature, `pctOverallCorrectSessionPM`, better represents the state in that it measures the student’s more recent performance. `PctOverallCorrectSessionPM` (Row two) occurred three times: one for NormGain and two for InvNormGain. The new features occurred a total of five times: “`nCorrectKCSessionPM** (2)`” in Row four, “`pctCorrectKCSessionPM** (1)`” in Row seven, and “`nIncorrectKCSessionPM** (2)`” in Row nine.

6.7.3.5 Background Features Much to the author’s surprise, only one background feature occurred in one final tutorial tactic: “`ageBG**`” (the age of the subject). The policy involved “`ageBG**`” is on KC_{24} and Justify/Skip-Justify. The remaining four background features were not involved in any policy.

6.7.3.6 Student Dialogue Features Student Dialogue Features are simple linguistic features that are computed from the student’s entries in the tutorial dialogue. Ten features were defined in this category. Table 6.11 summarizes the occurrences of each feature in the induced NormGain and InvNormGain tutorial tactics. The features in this category

Table 6.10: Occurrence of Performance Features in The Final Tutorial Tactics

		NormGain (5)	InvNormGain (12)
1	pctCorrectPM (1)	0	1
2	pctOverallCorrectSessionPM (3)	1	2
3	nCorrectKCPM (1)	0	1
4	nCorrectKCSessionPM** (2)	1	1
5	pctOverallCorrectKCPM (1)	0	1
6	pctCorrectKCPM (1)	0	1
7	pctCorrectKCSessionPM** (1)	0	1
8	nIncorrectKCPM (5)	2	3
9	nIncorrectKCSessionPM** (2)	1	1

occurred fourteen times. Among them, “stuAverageWordsSD**” (the average number of words per student turn) occurred four times, while stuConceptToWordRatioSD** (the ratio of physics concept words to total words in the student’s turns) occurred three times. Three features: averagePhysConceptsStudentDialogueSD**, stuAverageConceptSD**, averagePhysConceptsStudentDialogueSessionSD** did not appear in any of the final tutorial tactics.

To summarize, Problem Solving Contextual Features occurred most frequently, fifty-eight times, in the final thirty-four induced tutorial tactics. Background Features occurred the fewest number of times. The newly added features were involved in the final in a total of $2 + 41 + 5 + 1 + 14 = 63$ times, so it could be concluded that expanding feature choices to include new features such as *conceptDifficultyPS***, was a good decision given that it had the most occurrences in the final tutorial tactics.

Table 6.11: Occurrence of Student Dialogue Features in The Final Tutorial Tactics

		NormGain (8)	InvNormGain (6)
1	numStudentConceptualDialogueSD** (1)	1	0
2	stuConceptToWordRatioSD**(3)	1	2
3	stuAverageWordsSD** (4)	2	2
4	numStudentConceptualDialogueSessionSD** (1)	1	0
5	stuConceptToWordRatioSessionSD** (2)	1	1
6	stuAverageWordsSessionSD** (1)	0	1
7	stuAverageConceptSessionSD** (2)	2	0

6.7.4 Feature Selection

In this study, I applied 11 feature selection methods. It would be interesting to see which of them found the most final tutorial tactics. Table 6.12 lists all the feature selection methods that were followed to get the final tutorial tactics for the corresponding KCs on the two types of tutorial tactics: the NormGain and InvNormGain ones. Additionally, “single” means it is a single feature policy.

It can be concluded that the three feature selection approaches: PCA-only, PCA-ECR, and PCA-Upper_Bound did not elicit any of the final tutorial tactics. All other eight approaches resulted in at least one. Among them, the two RL-based feature selection methods appeared to be most effective. The ECR-based method discovered four NormGain tutorial tactics and six InvNormGain tutorial tactics. The Upper_Bound method found five NormGain tutorial tactics and four InvNormGain tutorial tactics. The feature selection may still need to be improved because one of the final induced policies is from the random feature selection — $\pi(KC_{20}, JS, InvNormGain)$.

Table 6.12: Applying 11 Feature Selection Methods to Induce 34 Tutorial Tactics

	NormGain		InvNormGain	
	ET	JS	ET	JS
KC_1	single		Upper_Bound	
KC_{14}	single	single	single	single
KC_{20}	ECR	PCA-Hedge	PCA-Lower_Bound	Random
KC_{21}	Upper_Bound	PCA-Hedge	Hedge	ECR
KC_{22}	Hedge	Upper_Bound	ECR	ECR
KC_{24}	ECR	Upper_Bound	ECR	Upper_Bound
KC_{27}	PCA-Random	ECR	Lower_Bound	Hedge
KC_{28}	Upper_Bound	Upper_Bound	ECR	Upper_Bound
Overall	Lower_Bound	ECR	Upper_Bound	ECR

6.8 SUMMARY: RL IN STUDY 3

To summarize, compared with the RL approach in Study 2, a series of changes were made in Study 3 to improve the effectiveness of the induced tutorial policies. However, there are many ways this can be explored in future work. For example, except for following the heuristic among the KCs to resolve the conflicting policies, we could choose the action that has the most votes. For instance, if three KCs, KC_{20} , KC_{21} and KC_{24} , were involved in an ET tutorial decision step. The ET policies on KC_{20} and KC_{24} selected to tell, but the KC_{21} 's policy selected to elicit, then the system would choose the one with the majority of the votes—tell in this case. Similarly, instead of using the Exploratory Corpus and DichGainDichGain Corpus individually or combined, a subset of the student's dialogue from each corpus could be selected to make a new training Corpus, and so on. Other research opportunities exist which will be discussed in Chapter 9.

The RL approach in Study 3 showed that it appears that the Problem Solving Contextual features are most involved in the final induced tutorial tactics and that most of the NormGain tutorial tactics were derived from the Exploratory Corpus while most of the InvNormGain tutorial tactics were derived from the DichGain Corpus. Among the 11 feature selection approaches, it seemed that the two RL-based tutorial tactics: Upper-Bound and ECR, were most effective. However, in order to investigate why these are the case, we need more exploration and it is beyond the content of this dissertation. Additionally, one of the future works I would like to investigate is how different choices of Training Corpus or feature selection methods are correlated with learning gains.

Next, the induced tutorial tactics were evaluated on real human subjects to see whether the students who followed the NormGain tutorial tactics would out-perform those who are under the InvNormGain ones.

7.0 STUDY 3: NORMALIZED GAIN (NORMGAIN) VS. INVERSE NORMALIZED GAIN (INVNORMGAIN)

In this chapter, I will present an experimental comparison of the induced Normalized Gain (NormGain) and Inverse Normalized Gain (InvNormGain) tutorial tactics.

7.1 METHODS

7.1.1 Participants

Data was collected over a period of two months during the summer of 2009. Participants were 64 college students who received payment for their participation. They were required to have a basic understanding of high-school algebra. However, they could not have taken any college-level physics courses. Students were randomly assigned to the two conditions. Each took from one to two weeks to complete the study over multiple sessions. In total, 57 students completed the study (29 in the NormGain condition and 28 in the InvNormGain condition).

7.1.2 NormGain-Cordillera and InvNormGain-Cordillera

NormGain-Cordillera and InvNormGain-Cordillera were used in Study 3. The only differences between the two systems were that interaction decisions made by NormGain-Cordillera were guided by the 17 NormGain tutorial tactics and those made by InvNormGain-Cordillera

were guided by the 17 InvNormGain tutorial tactics. Only one human wizard (the author) was involved in Study 3.

7.1.3 Materials & Procedures

Participants in both conditions experienced a background survey; read a textbook covering the target domain knowledge; took a pre-test; solved the same seven training problems in the same order on the NormGain-Cordillera for the NormGain condition and the InvNormGain-Cordillera for InvNormGain condition; and finally took a post-test that was identical to the pre-test. The exams, introductory materials and training problems used were identical to those used in Study 2 as was their order of presentation. As discussed in Chapter 5 subsection 5.2.2, in Study 3, everything but the exams were identical to those used in Study 1. The exams differed in only a single test question.

7.1.4 Grading

In study 3, the students' test answers were graded by a different grader, the author, than the grader from Studies 1 and 2. All of the tests in study 3 were graded in a double-blind manner and followed the same procedure as used by the grader in Studies 1 and 2. Following the same grading rubrics used in the previous studies, each question was assigned two grades: overall and KC-based grade.

A grader agreement study was conducted in order to establish the validity of this grading. The inter-grader agreement study, discussed in Appendix B, showed that the grading rubrics were comparable to those used in the prior studies.

7.1.5 Measures

The main purpose of Study 3 was to investigate whether micro-level tutorial decisions would make a difference in learning. The hypothesis was that the NormGain group would outperform the InvNormGain group. Therefore, Study 3 mainly focused on the two groups' learning performances, which is measured by their pre-test and post-test scores. Moreover,

the students' overall learning performance will be compared under both grading criteria.

7.2 RESULTS

Random assignment appears to have balanced the incoming student competence across conditions. A post-evaluation analysis showed that there were no statistically significant differences in the pre-test scores between the two conditions. Additionally, there were no significant differences between two groups on the mathSAT scores with a one-tailed paired t-test: $t(39) = 0.536, p = 0.595$ ($M = 633.48, SD = 140.14$ for the NormGain group and $M = 654.55, SD = 108.92$ for the InvNormGain group) or age: $t(55) = 0.175, p = 0.862$ ($M = 23.41, SD = 4.39$ for the NormGain group and $M = 23.64, SD = 5.47$ for the InvNormGain group).

7.2.1 Time

No significant difference was found between the two groups in terms of the total training time spent on Cordillera with a one-tailed paired t-test: $t(55) = -.272, p = .787$. The NormGain group spent ($M = 259.98, SD = 59.22$) and the InvNormGain group spent ($M = 264.57, SD = 67.60$). By using a one-tailed paired t-test, a more detailed analysis of the time spent on a per-problem basis revealed no significant difference between the two groups for all but P6. Figure 7.1 compares the average time students spent on each training problem between the two groups. On P6, the InvNormGain group spent an average of 57.72 min on the problem. This was significantly more than the NormGain group's 47.09 min ($t(55) = 3.28, p = 0.002$).

7.2.2 Learning Performance

7.2.2.1 Compare NormGain vs. InvNormGain conditions: Overall Learning Performance A one-way ANOVA was used to test for learning performance differences between the pre- and posttests. Participants in Study 3 made significant gains from pre-test

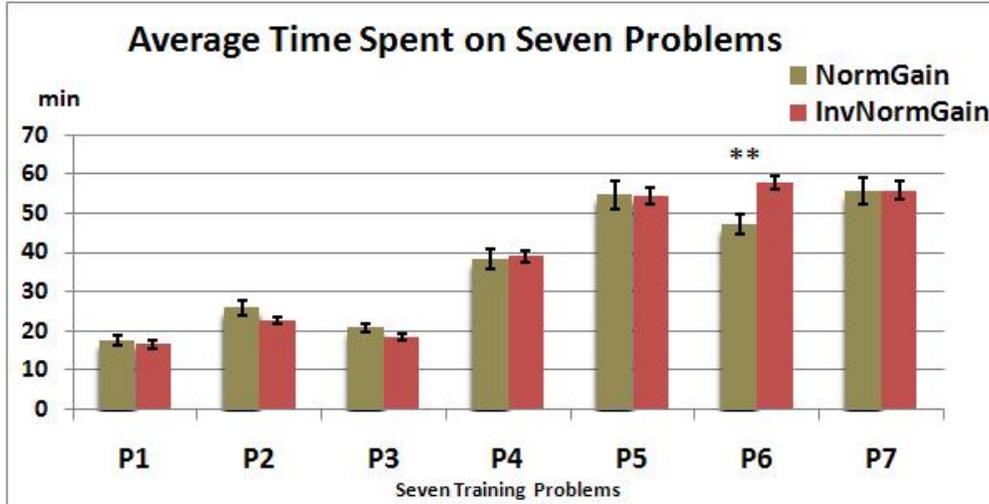


Figure 7.1: Compare Time Between NormGain vs InvNormGain Groups On Training Problem

to post-test: with $F(1, 112) = 36.22, p = .000, R^2 = .70$ under the overall grading criteria and $F(1, 112) = 27.58, p = .000, R^2 = 0.71$ under the cumulative KC-based grading criteria.

Table 7.1 compares the pre-test, post-test, adjusted-post-test, and NLG scores between the two conditions under two grading criteria. In Table 7.1, the Adj.Post-test scores were compared between the two conditions by running an ANCOVA using the corresponding pre-test score as the covariate. The rest of the three scores, the pre-test, post-test and NLG scores, were compared with one-tailed paired t-tests. The first column in Table 7.1 shows that there were two grading criteria and the second column shows the comparisons were conducted on four types of test scores: pre-test scores, post-test scores, adjusted post-test scores, and NLG under each grading criteria. The third and fourth columns in Table 7.1 list the means and SDs σ of the NormGain and InvNormGain groups' corresponding scores. The fifth column lists the corresponding statistical comparison and the sixth column lists the effect size of the comparison. Similar to Study 2, in study 3 Cohen's d , was used. This is defined as the mean learning gain of the experimental group minus the mean learning gain

of the control group, divided by the groups' pooled standard deviation. The final column lists the statistical power of the comparison, $1 - \beta$.

Table 7.1 shows that there was no significant difference between the two groups on pre-test scores under either grading criteria. However, there was a significant difference between the two groups on the post-test, adjusted-post-test, and NLG scores under both grading criteria. For example, cumulative KC-based NLG scores were used as reward functions for inducing NormGain and cumulative KC-based (1-NLG) scores were used as reward functions for inducing InvNormGain. The last row in Table 7.1 shows that with a one-tailed paired t-test the former group out-performed the latter group: $t(55) = 3.058, p = 0.003$ and the effect size was 0.81. Across all measurements, the NormGain group performed significantly better than the InvNormGain group and the effect size was large by Cohen's d criteria.

7.2.2.2 Compare NormGain vs. InvNormGain Conditions: KC-based Learning

Performance Above students' overall learning performance was compared in the previous section. Since KC-based tutorial tactics were induced, it would be interesting to compare the two groups' performance on a KC basis. With first step was to investigate whether students learned on the eight primary KCs. A one-way ANOVA was used to test for learning performance differences between the pre- and post-tests on each KC. Table 7.2 shows that the participants in study 3 gained significantly from pre-test to post-test on all eight primary KCs.

Next, the two groups' KC-based scores were compared. Table 7.3 compares the two groups' performance on the pre-test, post-test, adjusted-post-test, and NLG scores on the eight primary KCs. If there was a difference between the two groups, the Cohen d value is labeled with “**” and it is labeled with “*” if it is marginal significant ($p < 0.1$). No significant difference was found between the two groups on the KC-based pretest scores across all eight KCs. Only on KC_{27} , did the NormGain group score marginally higher than the InvNormGain group.

On four out of eight primary KCs (KC_1 , KC_{20} , KC_{21} , and KC_{27}) the NormGain group significantly out-performed the InvNormGain group on the post-test, adjusted post-test and NLG scores. On KC_{22} and KC_{24} , the NormGain group significantly out-performed the In-

Table 7.1: NormGain vs. InvNormGain on Pre- and Post-Test

	NormGain	InvNormGain	Stat	cohen d ^a	1 - β
Overall	Pretest	0.48 (0.18) ^b	$t(55) = 0.719, p = 0.475$	0.20	0.59
	Posttest	0.72 (0.11)	$t(55) = 2.817, p = 0.007$	0.75 **	0.51
	Adjusted Posttest	0.71(0.078)	$F(1, 54) = 18.33, p = 0.000$	1.17 **	0.81
	NLG	0.48 (0.16)	$t(55) = 3.764, p = 0.000$	0.99 **	0.59
Cumulative KC-based	Pretest	0.42 (0.16)	$t(55) = 0.71, p = .484$	0.15	0.55
	Posttest	0.65 (0.15)	$t(55) = 2.32, p = 0.024$	0.65 **	0.53
	Adjusted Posttest	.63 (.095)	$F(1, 54) = 10.689, p = .002$	0.86 **	0.49
	NLG	0.41 (0.19)	$t(55) = 3.058, p = 0.003$	0.81 **	0.48

^aIn this dissertation, I used Cohen's d, which is defined as the mean learning gain of the experimental group minus the mean learning gain of the control group, divided by the groups' pooled standard deviation.

^bStd. Deviation

vNormGain group on post-test and adjusted post-test scores, but not on the NLG scores. On KC_{14} , the NormGain group significantly out-performed the InvNormGain group on adjusted post-test score, but only marginally significantly on the post-test and NLG scores. On KC_{28} , however, no significant difference was found between the two groups on post-test, adjusted post-test or NLG scores. One of the potential explanations for the lack of difference on learning between the two conditions across measurements may be that there were too few times KC_{28} appeared during the training. For example, Table 4.2 showed that KC_{28} only appeared during the problem solving for training problem P4 while the other KC that has the equivalent complexity is KC_{27} (KC_{27} and KC_{28} are the only two top level KCs.), which showed up in three training problems: P5, P6 and P7.

Table 7.2: KC-based Pre- and Post-Test Test Scores

	Pretest	Posttest	Stat
KC_1	0.40 (0.19)	0.59 (0.19)	$F(1, 112) = 26.67, p < 0.0000, R^2 = 0.582$
KC_{14}	0.44 (0.23)	0.59 (0.25)	$F(1, 112) = 10.85, p = 0.001, R^2 = 0.46$
KC_{20}	0.37 (0.20)	0.62 (0.16)	$F(1, 112) = 54.27, p < 0.0000, R^2 = 0.48$
KC_{21}	0.44 (0.22)	0.70 (0.18)	$F(1, 112) = 46.37, p < 0.0000, R^2 = 0.42$
KC_{22}	0.40 (0.25)	0.57 (0.24)	$F(1, 112) = 12.78, p = 0.001, R^2 = 0.44$
KC_{24}	0.43 (0.19)	0.61 (0.18)	$F(1, 112) = 25.68, p < 0.0000, R^2 = 0.58$
KC_{27}	0.48 (0.23)	0.68 (0.26)	$F(1, 112) = 19.98, p < 0.0000, R^2 = 0.52$
KC_{28}	0.36 (0.23)	0.50 (0.23)	$F(1, 55) = 10.35, p = 0.002, R^2 = 0.48$

Table 7.3: Between-Group Comparison by KC-based
Pre- and Post-Test Scores

KC	Test	NormGain	InvNormGain	Stat	d	$1 - \beta$
KC_1	Pre	0.42 (0.15)	0.39 (0.22)	$t(55) = 0.66, p = 0.5095$	0.16	0.58
	Post	0.65 (0.16)	0.53 (0.21)	$t(55) = 2.51, p = 0.0151$	0.66 ^{**1}	0.48
	Adj.	0.64 (0.12)	0.54 (0.12)	$F(1, 54) = 9.80, p = 0.0028$	0.85 ^{**}	0.52
	NLG	0.41 (0.24)	0.24 (0.27)	$t(55) = 2.59, p = 0.0122$	0.68 ^{**}	0.48
KC_{14}	Pre	0.43 (0.23)	0.44 (0.25)	$t(55) = -0.17, p = 0.8638$	-0.04	0.87
	Post	0.64 (0.22)	0.53 (0.26)	$t(55) = 1.71, p = 0.0937$	0.47 ^{*2}	0.52
	Adj.	0.65 (0.17)	0.53 (0.17)	$F(1, 54) = 6.47, p = 0.0139$	0.72 ^{**}	0.56
	NLG	0.38 (0.34)	0.18 (0.39)	$t(53) = 2.00, p = 0.0502$	0.56 [*]	0.54
KC_{20}	Pre	0.38 (0.17)	0.37 (0.22)	$t(55) = 0.31, p = 0.7613$	0.05	0.77
	Post	0.68 (0.13)	0.57 (0.18)	$t(55) = 2.48, p = 0.0163$	0.72 ^{**}	0.58
	Adj.	0.67 (0.11)	0.58 (0.11)	$F(1, 54) = 10.30, p = 0.0022$	0.83 ^{**}	0.47
	NLG	0.47 (0.22)	0.32 (0.19)	$t(55) = 2.75, p = 0.0080$	0.74 ^{**}	0.51
KC_{21}	Pre	0.45 (0.20)	0.43 (0.24)	$t(55) = 0.35, p = 0.7256$	0.09	0.74
	Post	0.75 (0.12)	0.65 (0.21)	$t(55) = 2.32, p = 0.0238$	0.6 ^{**}	0.47
	Adj.	0.75 (0.13)	0.65 (0.13)	$F(1, 54) = 7.62, p = 0.0079$	0.78 ^{**}	0.57
	NLG	0.56 (0.22)	0.36 (0.32)	$t(55) = 2.73, p = 0.0086$	0.74 ^{**}	0.53
KC_{22}	Pre	0.42 (0.25)	0.39 (0.26)	$t(55) = 0.41, p = 0.6828$	0.12	0.71
	Post	0.64 (0.19)	0.50 (0.27)	$t(55) = 2.34, p = 0.0228$	0.61 ^{**}	0.48
	Adj.	0.63 (0.17)	0.51 (0.17)	$F(1, 54) = 7.77, p = 0.0073$	0.72 ^{**}	0.47
	NLG	0.32 (0.43)	0.22 (0.33)	$t(54) = 0.97, p = 0.3380$	0.26	0.54
KC_{24}	Pre	0.46 (0.15)	0.41 (0.23)	$t(55) = 0.89, p = 0.3782$	0.26	0.57
	Post	0.65 (0.14)	0.56 (0.20)	$t(55) = 2.03, p = 0.0468$	0.53 ^{**}	0.49
	Adj.	0.64 (0.11)	0.58 (0.11)	$F(1, 54) = 4.22, p = 0.0448$	0.56 ^{**}	0.51

Continued on Next Page...

¹ “**” means significant.

² “*” means marginal significant.

Table 7.3: Between-Group Comparison by KC-based Pre- and Post-Test Scores

KC	Test	NormGain	InvNormGain	Stat	d	$1 - \beta$
	NLG	0.37 (0.22)	0.27 (0.27)	$t(55) = 1.67, p = 0.1012$	0.41	0.46
KC_{27}	Pre	0.53 (0.21)	0.42 (0.24)	$t(55) = 1.74, p = 0.0879$	0.5*	0.55
	Post	0.78 (0.21)	0.58 (0.28)	$t(55) = 3.01, p = 0.0040$	0.82**	0.54
	Adj.	0.74 (0.18)	0.63 (0.18)	$F(1, 54) = 5.88, p = 0.0187$	0.62**	0.47
	NLG	0.60 (0.34)	0.28 (0.37)	$t(53) = 3.33, p = 0.0016$	0.92**	0.55
KC_{28}	Pre	0.37 (0.20)	0.36 (0.26)	$t(55) = 0.13, p = 0.8997$	0.04	0.9
	Post	0.53 (0.22)	0.47 (0.24)	$t(55) = 1.01, p = 0.3179$	0.27	0.52
	Adj.	0.53 (0.17)	0.47 (0.17)	$F(1, 54) = 1.61, p = 0.2101$	0.36	0.54
	NLG	0.27 (0.37)	0.16 (0.29)	$t(54) = 1.26, p = 0.2119$	0.34	0.51

7.2.2.3 Summary of Learning Results showed that both groups of participants had significant learning gains after training on NormGain-Cordillera and InvNormGain-Cordillera respectively. More importantly, although no significant difference was found in time on task and in the pre-test scores under both grading criteria, the NormGain group out-performed the InvNormGain group on the post-test, adjusted post-test, and NLG scores regardless of the grading criteria. Therefore overall, the results show that the micro-level tutorial decisions on micro-steps made a significant difference in the students' learning.

On a KC by KC basis, the difference between the two groups was not significant on all of the eight primary KCs. Especially, on KC_{28} , no significant difference was found between the two groups on post-test, adjusted post-test, and NLG scores. There are many potential explanations for the lack of the difference on KC_{28} . For example, KC_{28} is the only KC for which both ET and JS tutorial tactics were derived from the DichGain Corpus. One of the hypothesis generated from this study was that Exploratory Corpus seemingly works more effectively than the DichGain Corpus (this will be discussed in the next chapter).

To summarize, the overall test scores seemingly support the primary research hypothesis. The NormGain condition indeed out-performed the InvNormGain condition. In order to investigate why the NormGain tutorial tactics were more effective than the InvNormGain one, it will be necessary to dig into the logs to have a detailed comparison of the differences between the two sets of tutorial tactics. For example, the induced NormGain tutorial tactics might simply elicit more answers from the students or execute more justification steps during the tutoring. Therefore, the following section will investigate whether the NormGain and InvNormGain tutorial tactics resulted in a different number of tutorial actions. The number of overall decisions the tutor made, the number of ET decisions, the I-ratio, and the number of JS decisions and the J-ratio between the two groups will all be compared. The goal is to see whether the NormGain tutorial tactics resulted in different tutorial behaviors from the InvNormGain policies when viewed from this shallow aspect.

7.2.3 Log Analysis

Table 7.4: Overall Characteristics of Tutorial Decisions in Exploratory Corpus

		NormGain (29)	InvNormGain (28)	Stats
1	Tell	63.759 (19.528)	63.250 (4.656)	$t(55) = 0.134, p = 0.894$
2	Elicit	198.586 (17.463)	204.000 (7.679)	$t(55) = -1.506, p = 0.138$
3	ET Decisions	262.345 (6.149)	267.250 (6.775)	$t(55) = -2.864, p = 0.006$
4	Skip-Justify	9.345 (3.829)	11.000 (1.700)	$t(55) = -2.096, p = 0.041$
5	Justify	42.517 (3.786)	40.321 (1.442)	$t(55) = 2.874, p = 0.006$
6	JS Decisions	51.862 (0.833)	51.321 (1.156)	$t(55) = 2.030, p = 0.047$
7	Overall Decisions	280.103 (4.126)	285.464 (6.995)	$t(55) = -3.539, p = 0.001$
8	I-ratio	0.758 (0.073)	0.763 (0.018)	$t(55) = -0.395, p = 0.694$
9	J-ratio	0.820 (0.073)	0.786 (0.030)	$t(55) = 2.273, p = 0.027$

7.2.3.1 Overall Tutorial decision Steps Table 7.4 summarizes and compares the average number of various tutorial decisions, the I-ratio and J-ratio between the NormGain and InvNormGain tutorial dialogues. Row 1 to row 7 shows the average numbers of various tutorial decisions the students got during the tutoring. These include the average number of tell decisions (row 1), elicit decisions (row 2), ET decisions (row 3), skip-justify decisions (row 4), justify decisions (row 5), JS decisions (row 6), and overall decisions (row 7). Table 7.4, shows that except for the total number of tells and elicits, the two groups differed in all other five numbers.

On average, the InvNormGain-Cordillera made more tutorial decisions during the tutoring than the NormGain-Cordillera. This is probably because the InvNormGain students got more remediations in their dialogues. Row 3 shows that the InvNormGain students had more ET decisions in their tutorial dialogue than the NormGain students. For JS decisions, the NormGain students got more justification steps than the InvNormGain ones: on average two more justification decisions (row 5).

Rows 8 and 9 compared the I-ratio and J-ratio between the two conditions. There were no significant differences between the two groups on the I-ratio. However, on J-ratio, the NormGain students were higher than the InvNormGain group. So by following the induced tutorial tactics, the NormGain tutorial dialogues seemingly were no more interactive than the InvNormGain ones but the NormGain students were more likely to get justification steps.

7.2.3.2 Comparing I-ratio Across Primary KCs Although no significant difference was found between the two groups on the I-ratio overall, once the dialogue was broken into a KC by KC basis there were significant differences between the two groups on each of the eight primary KCs (see Table 7.5). In Table 7.5, row 2 shows that on KC_{14} the NormGain group got all elicits while the InvNormGain group got all tells. Among the rest of seven primary KCs, the NormGain condition was more likely to get elicits than the InvNormGain condition on KC_{20} , KC_{21} , and KC_{22} ; and the InvNormGain condition was more likely to get elicits than the NormGain condition on KC_1 , KC_{24} , KC_{27} , and KC_{28} .

Table 7.5: Compare NormGain vs. InvNormGain on I-ratio Across Eight Primary KCs

		<i>NormGain(29)</i>	<i>InvNormGain (28)</i>	<i>Stats</i>
1	KC_1	0.500 (0.000)	0.696 (0.157)	$t(55) = -6.72, p = 0.000$
2	KC_{14}	1.000 (0.000)	0.000 (0.000)	
3	KC_{20}	0.897 (0.024)	0.696 (0.030)	$t(55) = 27.87, p = 0.000$
4	KC_{21}	0.923 (0.030)	0.863 (0.045)	$t(55) = 5.95, p = 0.000$
5	KC_{22}	0.888 (0.099)	0.543 (0.089)	$t(55) = 13.88, p = 0.000$
6	KC_{24}	0.866 (0.028)	0.920 (0.029)	$t(55) = -7.21, p = 0.000$
7	KC_{27}	0.484 (0.137)	0.651 (0.112)	$t(55) = -5.03, p = 0.000$
8	KC_{28}	0.000 (0.000)	0.525 (0.108)	$t(55) = -26.08, p = 0.000$

7.2.3.3 Comparing J-ratio Across Primary KCs Similarly, the J-ratio was broken into a KC by KC basis. Only seven primary KCs (KC_1 was not involved in JS decisions) were involved in JS decisions (see Table 7.6). Surprisingly, on two KCs, KC_{22} (row 4) and KC_{28} (row 7), both NormGain and InvNormGain tutorial tactics achieved the same results, executing all justification steps. There are at least two potential explanations. One possible explanation is that the JS decisions on these KCs may not matter to the students' learning. The other possible explanation is that the source training corpora used to induce these two KC-specific policies might not be exploratory enough. On KC_{14} , however, following the NormGain tutorial tactics resulted in skipping all justification steps, but following the InvNormGain tutorial tactics resulted in executing all justification steps. For the remaining four KCs, no significant difference was found between the two conditions on KC_{20} (row 2) and KC_{24} (row 5). Only marginally significant difference was found between the two groups on KC_{27} . On KC_{21} , however, the NormGain group was significantly more likely to get justification steps than the InvNormGain group.

Table 7.6: Compare NormGain vs. InvNormGain on J-ratio across Eight Primary KCs

		NormGain(29)	InvNormGain (28)	Stats
1	KC_{14}	0.000 (0.000)	1.000 (0.000)	
2	KC_{20}	1.000 (0.000)	0.994 (0.022)	$t(55) = 1.467, p = 0.148$
3	KC_{21}	0.815 (0.216)	0.573 (0.096)	$t(55) = 5.445, p = 0.000$
4	KC_{22}	1.000 (0.000)	1.000 (0.000)	
5	KC_{24}	0.876 (0.024)	0.871 (0.005)	$t(55) = 1.071, p = 0.289$
6	KC_{27}	0.046 (0.140)	0.000 (0.000)	$t(55) = 1.736, p = 0.088$
7	KC_{28}	1.000 (0.000)	1.000 (0.000)	

7.2.3.4 Summary of Log Analysis Following the NormGain tutorial tactics did not generate more interactive tutorial tactics than following the InvNormGain ones. But once broken into a KC by KC basis, the NormGain tutorial tactics resulted in different I-ratio for each of the primary KCs. On the other hand, following the NormGain tutorial tactics seemed more likely to execute a justification step but once broken it into KC by KC bases, the NormGain and InvNormGain tutorial tactics' J-ratio were only significantly different on KC_{21} and KC_{14} (The NormGain tutorial tactics skipped all of them while the InvNormGain executed all of them). To summarize, future work is needed to explore why the NormGain tutorial policies resulted in better learning performance than the InvNormGain ones.

7.3 DISCUSSION

To summarize, the findings confirmed the primary hypotheses in this thesis. First, and foremost, the pedagogical tutorial tactics applied at the interactive decision level affected students' learning. Secondly, the use of RL to derive tutorial tactics from existing data proved to be feasible and successful. On the other hand, the results also suggested that content

exposure with the Cordillera system, irrespective of the tactics employed, was, indisputably, an important factor in governing students' learning even the InvNormGain students learned significantly in Study 3. Despite this importance, however the results showed that the pedagogical tutorial tactics also made a significant impact.

However, it is not clear as to what it was about the induced NormGain tutorial tactics that caused the NormGain students to learn more effectively than the InvNormGain group. By simply analyzing the log file in a relatively shallow way, it seems that it was not that the NormGain tutorial tactics were simply more interactive or generated more justification steps that caused the NormGain students to learn more than the InvNormGain students. Overall, the preliminary results supported the conjecture that interactivity is not, necessarily, the most important determiner of students learning. For example, no significant difference was found between the two conditions in terms of the number of elicitation prompts and tells they received and the I-ratio. However, NormGain students learned significantly more than the InvNormGain students. Additionally, once broken into a KC by KC basis, the InvNormGain students had significantly higher I-ratio than the NormGain group on KC_1 , KC_{24} , KC_{27} , and KC_{28} , but the former did not learn more than the NormGain group.

For JS decisions, following the induced NormGain tutorial tactics indeed resulted in more justification steps in students' tutorial dialogues. However, once the tutorial decisions were broken into a KC by KC basis, the two groups differed significantly only on KC_{21} and KC_{14} . Therefore, future work is needed to investigate the induced tutorial tactics and find out what actually caused these learning differences.

The NormGain and InvNormGain tutorial tactics in Study 3 were derived from the Exploratory and DichGain Corpora in Studies 1 and 2. Therefore, it is possible to draw some hypothesis from observations by running a post-hoc comparison among the four groups. A cross-study analysis comparing the three studies will be presented in Chapter 8.

8.0 GENERAL DISCUSSION AND CONCLUSIONS

This chapter contains a general discussion of the results including a post-hoc comparison across the study groups. I will then revisit the central research questions examining the data relevant to them and draw conclusions.

8.1 POST-HOC COMPARISON

8.1.1 STUDY VARIATIONS

A total of 158 participants used four versions of Cordillera as part of the three studies: The Exploratory Group contained 64 students who used Random-Cordillera (Study 1); the Dichotic Gain (DichGain) Group was comprised of a total of 37 students who used DichGain-Cordillera (Study 2); The Normalized Gain (NormGain) group included 29 students who used NormGain-Cordillera (Study 3); and the Inverse Normalized Gain (InvNormGain) group included 28 students who used InvNormGain-Cordillera (Study 3). All of the participants followed the same procedure; used the same preparatory materials and problems; and interacted with Cordillera with the identical GUI. They all completed a background survey; read a textbook covering the target domain knowledge; took a pre-test; solved the same seven training problems in the same order on Cordillera; and finally took a post-test. Only four salient differences existed across the three studies:

1. Although all of the participants were recruited in the same way, they were recruited in different years. In Study 3 the students were randomly assigned into the NormGain and InvNormGain groups (2009). On the other hand, in the first two studies participants

were not randomly assigned to the Exploratory (2007) and DichGain groups (2008).

2. Interaction decisions that were made by NormGain-Cordillera, InvNormGain-Cordillera, DichGain-Cordillera, and Random-Cordillera were guided by different tutorial tactics. Random-Cordillera made random decisions on tutorial decision steps. The other three versions of Cordillera followed corresponding induced tutorial tactics to decide which action to take.
3. Apart from a single question variation on Studies 2 and 3, all three studies used identical exams containing a total of 33 test questions. The one variation occurred as the result of the replacement of a single question, Q_{20} , which had been used in Study 1. It was judged to be too easy and was replaced with a more difficult question, Q_{20}^* that covered the same KCs for Studies 2 and 3. The remaining 32 test items were identical across all three studies.
4. A group of six human wizards (including the author) were involved in Studies 1 and 2; but only one wizard (the author) was in Study 3.

Despite these differences, because the NormGain and InvNormGain groups trained in Study 3 were guided using tutoring tactics derived from the Exploratory and DichGain corpora, a post-hoc comparison among the four groups will allow us to observe the characteristics of the induced tutorial tactics from a wider point of view.

In Study 3, the exams were graded by a different grader, the author, than the grader in Studies 1 and 2. An inter-grader agreement study, discussed in Appendix B, showed that the grading rubrics were comparable to those used in the prior studies. Moreover, the high level of correlation supports the conclusion that the grades assigned by the two graders were equivalent and thus may be reliably compared. This chapter contains a post-hoc comparison across the four student groups. This analysis will make use of the new grades exclusively.

In order to establish test equivalence, Q_{20} and Q_{20}^* were excluded from the scores used here. As described in the previous chapter, the tests contained 33 test items which covered 168 KC occurrences. Removing Q_{20} reduced this total by 1 leaving 32 test items covering 166 KC occurrences. The subsections below compare learning on both the overall and cumulative KC-based scores. For the overall scores, the maximum raw score is 32 points while for the cumulative KC-based score it is 166. For comparison purposes both scores were normalized

to 1.

Based on the procedure of induced tutorial tactics, it was expected that $NormGain > DichGain > Exploratory > InvNormGain$. However, in Chapter 5, a post-hoc comparison showed no significant difference between the DichGain and Exploratory. These will be compared again using the new grading criteria.

A one-way ANOVA showed that there were no significant differences among the four groups on overall training time: ($F(3, 147) = 1.531, p = .209$). More specifically, the average total training time across the seven training problems, was $M = 278.73, SD = 67.38$ for Exploratory group, $M = 294.33, SD = 87.51$ for DichGain group, $M = 259.99, SD = 59.22$ for NormGain group, and $M = 264.57, SD = 67.60$ for InvNormGain group. Additionally, no significant difference was found among the Exploratory, the NormGain, and the InvNormGain groups on the MathSat scores: ($F(2, 83) = .520, p = .596$).

8.1.2 LEARNING PERFORMANCE

A one-way ANOVA was used to test for performance preference differences between the pre- and post-tests. Participants across four groups made significant gains from pre-test to post-test: $F(1, 314) = 67.36, p = .000, R^2 = .68$ under the overall grading criteria and $F(1, 314) = 41.82, p = .000, R^2 = 0.69$ under the cumulative KC-based grading criteria.

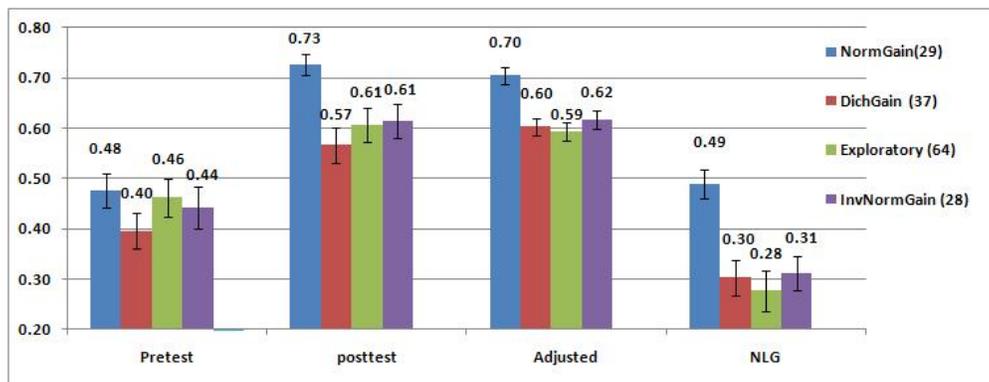


Figure 8.1: Compare Four Groups Learning Performance under Overall Grading

Figure 8.1 compares the four groups on the pre-test, post-test, adjusted-post-test, and NLG scores under the overall grading. A one-way ANOVA was used for learning performance differences among the four groups. No significant pre-test score differences were found between the groups under the overall-grading rubric ($F(3, 154) = 1.16, p = 0.32$). However, there were significant differences among the four groups on the remaining three scores: $F(3, 154) = 5.052, p = .002$ on posttest-scores, $F(3, 153) = 9.938, p = .000$ for adjusted post-test scores, and $F(3, 154) = 8.33, p = 0.000$ on the NLG scores.

Moreover, pairwise comparisons among the four groups¹ showed that there was a significant difference between the NormGain and either of the three groups on the post-test scores, adjusted post-test scores, and NLG scores. There were no significant differences among the DichGain, Exploratory, InvNormGain on all four test scores (see Table 8.1). More specifically, Table 8.1 shows the pairwise comparisons among the four groups on pre-test, post-test, adjusted post-test scores, and NLG scores. The first column lists the two groups in comparison and their corresponding mean and SD scores. The second column lists the statistical result of the t-test comparison. The last two columns list the effect size and power of the comparison. For effect size, Cohen's d was still used. Results in Table 8.1 suggest that under the overall grading rubric, on the pretest there were no significant differences among the four groups but $NormGain > DichGain = Exploratory = InvNormGain$ across post-test, adjusted post-test, and NLG scores. Although this formulation holds that the latter three groups were equally effective, in some tests the power was less than 0.80. So more subjects would be needed to determine which pairs of groups were in fact equivalent.

¹A Bonferroni correction was not performed as the hypotheses being tested are independent and the corpora were collected separately.

Table 8.1: Compare Four Groups Under the Overall Grading Criteria

<i>Group Name</i> $\mu(\sigma)$		<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
Pre-test				
NormGain	0.48(0.18)	$t(55) = 0.63, p = 0.53$	0.2	0.63
InvNormGain	0.44(0.22)			
NormGain	0.48(0.18)	$t(64) = 1.75, p = 0.09$	0.44	0.47
DichGain	0.40(0.19)			
NormGain	0.48(0.18)	$t(91) = 0.31, p = 0.75$	0.1	0.77
Exploratory	0.46(0.20)			
InvNormGain	0.44(0.22)	$t(63) = 0.91, p = 0.37$	0.2	0.49
DichGain	0.40(0.19)			
InvNormGain	0.44(0.22)	$t(90) = -0.43, p = 0.67$	-0.1	0.69
Exploratory	0.46(0.20)			
DichGain	0.40(0.19)	$t(99) = -1.65, p = 0.102$	-0.31	0.38
Exploratory	0.46(0.20)			
Posttest				
NormGain	0.73(0.11)	$t(55) = 2.85, p = 0.006$	0.77	0.51
InvNormGain	0.61(0.18)			
NormGain	0.73(0.11)	$t(64) = 4.12, p < 0.001$	1.04	0.42
DichGain	0.57(0.19)			
NormGain	0.73(0.11)	$t(91) = 3.28, p = 0.001$	0.74	0.27
Exploratory	0.61(0.18)			
InvNormGain	0.61(0.18)	$t(63) = 1.05, p = 0.30$	0.27	0.5
DichGain	0.57(0.19)			
InvNormGain	0.61(0.18)	$t(90) = 0.208, p = 0.836$	0.05	0.84
Exploratory	0.61(0.18)			
Continued on Next Page...				

Table 8.1: Compare Four Groups Under the Overall Grading Criteria

<i>Group Name</i>	$\mu(\sigma)$	<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
DichGain	0.57(0.19)	$t(99) = -1.04, p = 0.301$	-0.22	0.48
Exploratory	0.61(0.18)			
Adjusted Posttest				
NormGain	0.70(0.07)	$t(55) = 4.21, p < 0.0001$	1.14	0.52
InvNormGain	0.62(0.09)			
NormGain	0.70(0.07)	$t(64) = 4.70, p = 0.00001$	1.18	0.42
DichGain	0.60(0.10)			
NormGain	0.70(0.07)	$t(91) = 5.33, p = 0.00001$	1.21	0.2
Exploratory	0.59(0.10)			
InvNormGain	0.62(0.09)	$t(63) = 0.60, p = 0.551$	0.15	0.61
DichGain	0.60(0.10)			
InvNormGain	0.62(0.09)	$t(90) = 1.06, p = 0.293$	0.24	0.47
Exploratory	0.59(0.10)			
DichGain	0.60(0.10)	$t(99) = 0.45, p < 0.66$	0.09	0.68
Exploratory	0.59(0.10)			
NLG				
NormGain	0.49(0.16)	$t(55) = 3.95, p < 0.001$	1.06	0.52
InvNormGain	0.31(0.18)			
NormGain	0.49(0.16)	$t(64) = 4.35, p = 0.000$	1.09	0.43
DichGain	0.30(0.18)			
NormGain	0.49(0.16)	$t(91) = 4.67, p = 0.0000$	1.06	0.23
Exploratory	0.28(0.22)			
InvNormGain	0.31(0.18)	$t(64) = 0.199, p = 0.84$	0.05	0.85
DichGain	0.30(0.18)			
Continued on Next Page...				

Table 8.1: Compare Four Groups Under the Overall Grading Criteria

<i>Group Name</i>	$\mu(\sigma)$	<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
InvNormGain	0.31(0.18)	$t(90) = 0.733, p = 0.466$	0.17	0.55
Exploratory	0.28(0.22)			
DichGain	0.30(0.18)	$t(99) = -0.60, p < 0.552$	0.12	0.61
Exploratory	0.28(0.22)			

Under the Cumulative KC-based grading criteria, similar results were found. No significant pre-test score difference was found among the four groups under the Cumulative KC-based grading rubric ($F(3, 154) = 0.38, p = 0.77$). However, there were significant differences among the four groups on the remaining three scores: $F(3, 154) = 3.41, p = .02$ on post-test scores, $F(3, 153) = 8.09, p = .000$ for adjusted post-test scores, and $F(3, 154) = 5.30, p = 0.002$ on the NLG scores (see Table 8.2). Similar to overall grading criteria, t-test comparisons showed that there was a significant difference between the NormGain group and either of the three remaining groups on the post-test scores, adjusted post-test scores, and NLG scores and there were no significant differences among the DichGain, Exploratory, InvNormGain on all four test scores (see in Table 8.2). Therefore, the cumulative KC-based rubric results suggest that: $NormGain > DichGain = Exploratory = InvNormGain$ across all three different performance metrics: post-test, adjusted post-test, and NLG scores. Again, note that Table 8.2 shows that although this formulation holds that the latter three conditions were equally effective, in some tests the power was less than 0.80. So more participants would be needed to determine which pairs of groups were equally effective.

Table 8.2: Compare Four Groups Under the Cumulative
KC-based Grading Criteria

<i>Group Name</i> $\mu(\sigma)$		<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
Pretest				
NormGain	0.42(0.15)	$t(55) = 0.66, p = 0.507$	0.16	0.58
InvNormGain	0.39(0.23)			
NormGain	0.42(0.15)	$t(64) = 1.05, p = 0.299$	0.25	0.49
DichGain	0.38(0.17)			
NormGain	0.42(0.15)	$t(91) = 0.29, p = 0.792$	0.05	0.8
Exploratory	0.41(0.20)			
InvNormGain	0.39(0.23)	$t(63) = 0.16, p = 0.871$	0.05	0.87
DichGain	0.38(0.17)			
InvNormGain	0.39(0.23)	$t(90) = -0.5, p = 0.618$	-0.1	0.64
Exploratory	0.41(0.20)			
DichGain	0.38(0.17)	$t(99) = -0.81, p = 0.418$	-0.16	0.52
Exploratory	0.41(0.20)			
Posttest				
NormGain	0.65(0.15)	$t(55) = 2.32, p = 0.024$	0.64	0.53
InvNormGain	0.54(0.20)			
NormGain	0.65(0.15)	$t(64) = 3.28, p = 0.0017$	0.82	0.46
DichGain	0.50(0.21)			
NormGain	0.65(0.15)	$t(91) = 3.17, p = 0.0069$	0.63	0.35
Exploratory	0.53(0.21)			
InvNormGain	0.54(0.20)	$t(63) = 0.78, p = 0.439$	0.2	0.55
DichGain	0.50(0.21)			
InvNormGain	0.54(0.20)	$t(90) = 0.23, p = 0.820$	0.05	0.82
Exploratory	0.53(0.21)			
Continued on Next Page...				

Table 8.2: Compare Four Groups Under the Cumulative
KC-based Grading Criteria

<i>Group Name</i> $\mu(\sigma)$		<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
DichGain	0.50(0.21)	$t(99) = -0.68, p = 0.498$	-0.14	0.57
Exploratory	0.53(0.21)			
Adjusted Posttest				
NormGain	0.63(0.07)	$t(55) = 3.16, p = 0.003$	0.89	0.59
InvNormGain	0.55(0.11)			
NormGain	0.63(0.07)	$t(64) = 3.68, p = 0.0000$	1.18	0.38
DichGain	0.52(0.11)			
NormGain	0.63(0.07)	$t(91) = 4.77, p = 0.0000$	1.12	0.3
Exploratory	0.52(0.11)			
InvNormGain	0.55(0.11)	$t(63) = 1.20, p = 0.236$	0.28	0.45
DichGain	0.52(0.11)			
InvNormGain	0.55(0.11)	$t(90) = 1.24, p = 0.217$	0.28	0.43
Exploratory	0.52(0.11)			
DichGain	0.52(0.11)	$t(99) = -0.07, p = 0.945$	0	0.95
Exploratory	0.52(0.11)			
NLG				
NormGain	0.42(0.19)	$t(55) = 3.15, p = 0.0026$	0.87	0.54
InvNormGain	0.25(0.21)			
NormGain	0.42(0.19)	$t(64) = 4.626, p = 0.000$	0.95	0.18
DichGain	0.22(0.23)			
NormGain	0.42(0.19)	$t(91) = 3.61, p = 0.0005$	0.84	0.33
Exploratory	0.22(0.26)			
InvNormGain	0.25(0.21)	$t(63) = 0.546, p = 0.587$	0.14	0.63
DichGain	0.22(0.23)			
Continued on Next Page...				

Table 8.2: Compare Four Groups Under the Cumulative KC-based Grading Criteria

<i>Group Name</i>	$\mu(\sigma)$	<i>Stat</i>	<i>Cohen's d</i>	$1 - \beta$
InvNormGain	0.25(0.21)	$t(91) = 0.506, p = 0.61$	0.12	0.65
Exploratory	0.22(0.26)			
DichGain	0.22(0.23)	$t(55) = -0.045, p = 0.964$	0	0.96
Exploratory	0.22(0.26)			

8.1.3 LEARNING PERFORMANCE ACROSS THE FOUR GROUPS

To summarize, a post-hoc comparison of learning performance across three studies shows that the NormGain group significantly outperformed all other three groups while no difference was found between the remaining three groups. These results were consistent both for the adjusted post-test scores and the normalized learning gains. These results support the prior analysis of Study 3 which showed that the NormGain tutorial tactics significantly improved students' learning compared with the InvNormGain ones.

However, the lack of a significant difference between the InvNormGain, DichGain, and Exploratory groups seemingly contradicts the initial predictions. The InvNormGain strategies were specifically induced to enhance those decisions that contribute less or even none to the students' learning. Therefore, a lower performance on the students' part there than in at least the DichGain group, which sought to enhance the tutorial decisions that contribute to the students' learning, was expected. One possible explanation for the lack of difference is that the tutorial tactics employed by the DichGain- and Random-Cordillera systems were ineffective and thus presented a minimum bar. By 'ineffective' it does not mean that they prevented the students from learning but rather that they were not able to make a positive impact on their learning above and beyond the baseline provided by Cordillera itself. Here the basic practices and problems, domain exposure, and interactivity of Cordillera set a minimum bar of students' learning that the tactics, however poor, cannot prevent. This is only a

post-hoc explanation not a tested hypothesis, however it merits further study. On the other hand, note that under both grading criteria in some tests, the power of the comparisons among the three groups was less than 0.80. Therefore, more participants would be needed to determine which pairs of the three groups were truly equivalent in further studies.

8.1.4 LOG ANALYSIS

Having compared the individual groups' learning performance, this subsection will compare the log file variations across the four groups. Two types of tutorial actions are of principal interest in this dissertation: elicit/tell (ET) and justify/skip-justify (JS). Therefore, the following section will focus on two aspects of the ET and JS actions. The first aspect of interest is the *I-ratio* of the tutorial dialogue. That is, how often were the students given elicitation prompts in the course of the dialogue? The second is the *J-ratio*. This is the number of times the dialogue manager did not skip a justification step. Both values range from 0 (no elicits or justifies) to 1 (all elicitation or justification).

8.1.4.1 I-Ratio Table 8.3 summarizes t-test comparisons on the I-ratio among the four tutorial corpora. In Table 8.3, the first two columns list the two groups in comparison and their corresponding mean and SD scores. The last column lists the statistical results of the t-test comparisons. From the Table 8.3, the I-ratios for the four student groups were: 0.76 (NormGain), 0.76 (InvNormGain), 0.44 (DichGain), and 0.50 (Exploratory) respectively. Except for no significant difference between the NormGain and InvNormGain on the I-ratio, both groups were significantly more interactive than either the DichGain group or Exploratory group. Altogether, the result is $NormGain = InvNormGain > Exploratory > DichGain$ on the I-ratio.

Although high interactivity is a key characteristic of one-on-one human tutoring, the more successful tutorial tactics were not necessarily more interactive than the less successful tactics. Comparisons between the NormGain and InvNormGain groups suggest that it is not the absolute level of interactivity that determines the students' success. The NormGain group was more successful than the others despite there being no significant difference in

interactivity ratios between it and the InvNormGain group. Conversely, the InvNormGain group was no more successful than the Exploratory and DichGain groups despite being more interactive than either.

8.1.4.2 Justify Ratio Table 8.4 summarizes t-test comparisons on J-ratio among the four tutorial corpora. In Table 8.4, the first two columns list the two groups in comparison and their corresponding mean and SD scores. The last column lists the statistical results of the t-test comparisons. Table 8.4, shows that the mean of J-ratios for the four student groups were: 0.82 (NormGain), 0.79 (InvNormGain), 0.43 (DichGain), and 0.53 (Exploratory). The difference was statistically significant: $F(3, 154) = 322.88, p = 0.000$. Table 8.4 presents the pair wise t-test comparisons. It shows that on J-ratio, the result is: $NormGain > InvNormGain > Exploratory > DichGain$.

Table 8.5 summarizes and compares the average number of justification decisions students experienced during the tutoring among the four groups. The average number of steps for each student was: 42.52 for the NormGain group, 40.32 for the InvNormGain group, 24.89 for the DichGain group and 30.17 for the Exploratory group. The pairwise comparisons among the four groups are listed in Table 8.5. It shows that the number of justification steps students received were in the same order as the justification ratio: $NormGain > InvNormGain >$

Table 8.3: Pairwise Comparison Among Four Groups On I-ratio

<i>Group 1</i>		<i>Group 2</i>		<i>Group 1 vs. Group 2</i>
NormGain	0.76 (0.07)	InvNormGain	0.76 (0.02)	$t(55) = .395, p = .694$
NormGain	0.76 (0.07)	Exploratory	0.50 (0.03)	$t(91) = 24.72, p = 0.000$
NormGain	0.76 (0.07)	DichGain	0.44 (0.04)	$t(64) = 22.08, p = 0.000$
InvNormGain	0.76 (0.02)	Exploratory	0.50 (0.03)	$t(90) = 43.998, p = .000$
InvNormGain	0.76 (0.02)	DichGain	0.44 (0.04)	$t(63) = 36.34, p = .000$
Exploratory	0.50 (0.03)	DichGain	0.44 (0.04)	$t(99) = 7.967, p = .000$

Table 8.4: Pairwise Comparison Among Four Groups On J-ratio

<i>Group 1</i>		<i>Group 2</i>		<i>Group 1 vs. Group 2</i>
NormGain	0.82 (0.07)	InvNormGain	0.79 (0.03)	$t(55) = 2.27, p = .027$
NormGain	0.82 (0.07)	Exploratory	0.53 (0.06)	$t(91) = 18.95, p = 0.000$
NormGain	0.82 (0.07)	DichGain	0.43 (0.07)	$t(64) = 22.85, p = .000$
InvNormGain	0.79 (0.03)	Exploratory	0.53 (0.06)	$t(90) = 43.998, p = .000$
InvNormGain	0.79 (0.03)	DichGain	0.43 (0.07)	$t(63) = 26.65, p = .000$
Exploratory	0.53 (0.06)	DichGain	0.43 (0.07)	$t(99) = 7.894, p = .000$

Exploratory > *DichGain*.

Table 8.5: Pairwise Comparison Among Four Groups On Number of Justification Steps

<i>Group 1</i>		<i>Group 2</i>		<i>Group 1 vs. Group 2</i>
NormGain	42.52 (3.79)	InvNormGain	40.32 (1.44)	$t(55) = 2.87, p = 0.006$
		Exploratory	30.17 (3.83)	$t(91) = 19.33, p = 0.000$
		DichGain	24.89 (3.59)	$t(64) = 22.85, p = .000$
InvNormGain	40.32 (1.44)	Exploratory	30.17 (3.83)	$t(90) = 13.57, p = .000$
		DichGain	24.89 (3.59)	$t(63) = 21.45, p = .000$
Exploratory	30.17 (3.83)	DichGain	24.89 (3.59)	$t(99) = 6.83, p = .000$

To summarize, applying RL to induce tutorial tactics designed to enhance students' learning resulted in a set of tutorial tactics that involved substantially more justifications than tutorial tactics designed with the goal of enhancing those decisions that contribute less or even none to the students' learning. However, although the NormGain group had a higher ratio of justification prompts than the InvNormGain, Exploratory, or DichGain groups it is not the case that the absolute justification ratio guarantees learning. As with the interactivity ratio, the InvNormGain group received a higher justification ratio than the

Exploratory or DichGain groups despite having been induced to enhance those decisions that contribute less or even none to the students learning, and despite the absence of a significant difference in adjusted post-test scores or NLG between the groups. This is supported by the absolute number of justification steps as the InvNormGain group received more than ten more justification prompts on average than the Exploratory or DichGain groups while the NormGain group had only two more justification steps than the InvNormGain group. Given that the tutorial decisions in this dissertation were very fine-grained size level, an argument can be made something other than more justification steps causes more learning.

8.2 REVISITING THE TWO RESEARCH QUESTIONS

8.2.1 Question 1: Micro-level Pedagogical Tutorial Decisions Affect Students' Learning.

A comparison between the NormGain and InvNormGain groups from Study 3, discussed in Chapter 7, shows that tutorial tactics covering micro-level interaction decisions do affect students' learning. More specifically, the results support the hypothesis that interactive tutorial decisions such as the elicit/tell and justify/skip-justify decisions affect students' learning. In Study 3, the students were randomly assigned to balanced conditions and received identical training materials and procedures apart from the tutoring tactics employed. After spending the same amount of time on training, the NormGain group outperformed the InvNormGain group in terms of posttest scores, the adjusted post-test scores and the normalized learning gain regardless of the grading criteria.

In order to investigate why the NormGain tutorial tactics were more effective, some preliminary analyses of the students' log files was performed. It showed that the induced tutorial tactics did not produce more or less interactive tutorial dialogues for the NormGain group relative to the InvNormGain group. However the interactivity ratio (I-ratio) varied across the eight primary KCs. On some KCs, the NormGain group was more likely to be elicited for information rather than told it; while on other KCs, the NormGain group received more

didactic instruction. A wider comparison across the NormGain, InvNormGain, DichGain, and Exploratory groups suggests that it might not be increased interactivity that caused the NormGain students to learn more than the remaining three groups. Similarly, for justification decisions, although the NormGain group was given significantly more justification steps than the InvNormGain group overall, once the totals are broken down by KC, the two groups only differed on one KC, KC_{21} . Additionally, a wider comparison among the NormGain, InvNormGain, DichGain, and Exploratory groups suggests that increased learning might not be due to receiving a higher number of justification steps. The InvNormGain students had significantly more justification steps than the DichGain, and Exploratory groups. However, the former did not learn more than the latter two groups. Therefore, further analysis is needed to understand what caused the NormGain tutorial tactics to be more effective.

The analysis in Chapters: 4; 5; and 7 show that all four groups learned significantly by training on Cordillera. This result indicates that the content exposure and practice opportunities can cause students to learn even from tutors with poor pedagogical tutorial tactics. However, it also indicates that, with effective tutorial tactics, students can learn more and more effectively than without.

8.2.2 Question 2: Reinforcement Learning is a Feasible Method to Induce Tutorial Tactics.

The results so far suggest that NormGain-Cordillera outperformed InvNormGain-Cordillera system and further that it seems to be superior to the DichGain-Cordillera and Random-Cordillera as well. This success supports the hypothesis that RL-induced rules are effective and that the approach taken in Study 3 was a feasible one. However, inducing effective tutorial tactics is not trivial. In Study 2, the DichGain tutorial tactics did not seem to be more effective than the random decisions in Random-Cordillera. A number of factors were changed between Study 2 and Study 3. These included the choice of training corpora, the selections of knowledge components, the reward function, the feature choices, the maximum number of the features in a policy, the discretization procedure, the feature selection methods, and how the dialogue manager should respond when there were conflicting policies. All of

these factors might change the effectiveness of the resulting tactics. However, it is still not clear which factor or factors caused a change in effectiveness. Despite that, the results demonstrate the feasibility of RL to induce tutorial tactics. Or more accurately, the results show that: applying RL with a suitable training corpus; defining reward functions; and using the proper state representations on properly selected knowledge components; will produce effective tutorial tactics with a reasonable policy conflict heuristic.

Moreover, the RL-induced tutorial tactics induced in this dissertation seemingly to be highly adaptive both to the problem solving context and to other features. This fits in with the conjecture proposed by both learning and cognitive scientists that pedagogical skills and tutorial interactions should be adaptive to the tutorial context and students' needs. However more work remains to be done in this area to investigate whether it was the adaptiveness that caused the NormGain students to learn more effectively. In inducing the tutorial tactics the Expected Cumulative Reward (ECR) was used as the estimate of success. However both the InvNormGain and DichGain systems contained tactics whose ECR estimates far exceeded their actual performance with the students. More studies are necessary to identify good performance metrics that will allow researchers to evaluate candidate tutorial tactics without running expensive empirical studies.

9.0 CONTRIBUTIONS AND FUTURE WORK

In this dissertation Reinforcement Learning was applied to induce several sets of pedagogical tutorial tactics from existing tutoring corpora and then integrated them into a system for evaluation with human subjects. This study was mainly designed to: 1) examine the pedagogical importance of low-level interactive decisions in tutoring; and 2) test the viability of using reinforcement learning to induce pedagogical tutorial tactics. As such this is an interdisciplinary study that contributes to several fields.

In the field of cognitive science, this dissertation demonstrates that pedagogical skills governing low-level interactive tutorial decisions can impact students' learning. In the field of the learning science, it informs the ongoing discussion of interactive vs. didactic tutoring by suggesting that a tutor's success is not governed by *how often* they give interactive prompts or ask the students questions but *how well*. This dissertation also demonstrates that RL may be fruitfully applied to derive adaptive pedagogical tutorial tactics from student-computer interactivity data, thus informing the general field of AI and Education. Further this work demonstrates that existing a-priori theories about the importance of given features to tutoring may be assessed by means of induction and feature selection which seek to distinguish profitable feature choices from unprofitable ones thus informing the nascent field of Educational Data Mining.

9.1 CONTRIBUTION TO COGNITIVE & LEARNING SCIENCE

Examinations of pedagogical skills have long been a focus of attention for the instructional and learning sciences. For example, one preferred explanation for the effectiveness

of human one-on-one tutoring is that human tutors possess effective pedagogical skills that are responsible for the students' learning gains [Chi et al., 2001, Collins and Stevens, 1982, McArthur et al., 1982, Merrill et al., 1992]. Pedagogical skills generally involve the tutor's skillful execution of tutoring tactics such as eliciting student knowledge with a scaffolding question. While it is generally assumed that the tutor's interactive decisions are responsible for tutoring effectiveness [Collins and Stevens, 1982], little evidence has been presented to date demonstrating that either the human tutor has effective pedagogical skills or that pedagogical skills cause students to learn. In order to execute the pedagogical skills effectively, it is assumed that tutors should adapt their behaviors to the students' needs based upon their current knowledge level, general aptitude, emotional state and other salient features. Previous research, however has cast doubt on the tutor's ability to monitor the student's state accurately [Chi et al., 2004] and on whether they really adapt their tutorial decisions based on the present context. Chi, Siler, and Jeong for example, found that human tutors do not seem to maintain an accurate model of student's knowledge level during the tutoring process. Similarly, [Putnam, 1987] found that experienced tutors did not attempt to form detailed models of the students' knowledge before attempting remedial instruction. Rather, each teacher appeared to move through a general curricular script irrespective of the student's state. In recent years, some cognitive scientists have begun to doubt the effect of pedagogical skills on student's learning [Chi et al., 2001, Chi et al., 2004, VanLehn, 2006].

This dissertation investigated on pedagogical skills at a micro-step level. i.e. pedagogical tutorial tactics. These tactics do not govern the domain solution path selected for presentation or the problems presented. They only govern low-level tutorial interactions, e.g. whether the student is told what principle to apply or if the system elicits it from them with a prompt, and whether a student, once he/she has made a step, is asked to justify his/her answer or not. If fine-grained pedagogical skills of this type turn out to be effective, then more complex or content-oriented tactics, such as problem or sub-problem selection may be similarly effective. Both the Normalized Gain (NormGain) and the Inverse Normalized Gain (InvNormGain) groups in Study 3 solved the same training problems in the same training order, following the same solution path in the process by using the same tutorial scripts. The only difference between the groups was the tutorial tactics employed. Additionally, the Exploratory

and DichGain groups, though run at different times from the NormGain and InvNormGain groups, followed the same tutoring cycle save for the pedagogical tutorial tactics employed. And, like the NormGain and InvNormGain groups, the Exploratory and DichGain groups showed no salient difference in incoming competence. As reported in Chapters 6 and 8, the NormGain group outperformed the other three groups under either grading rubrics. Thus, a first, and primary, contribution of this work has been to show that pedagogical tutorial tactics can impact student’s learning. However, inducing such tutorial tactics is not trivial. As demonstrated in Studies 2 and 3, the induced DichGain tutorial tactics did not notably improve the effectiveness of Cordillera compared with the Random-Cordillera. Given the potentially important role tutorial tactics might play for the effectiveness of the tutoring, more research needs to be done to investigate this issue.

On the other hand, high interactivity is a key characteristic of one-on-one tutoring. A classroom lecture can be viewed as a monologue consisting of a long sequence of tutor instructions or “tell” acts. Individual tutoring, by contrast, features a great deal of give and take and can be viewed as a mixture of tutor questions or elicitation acts, student responses, and tutor instructions. One common assumption, often referred as the monotonic interaction hypothesis [VanLehn et al., 2007a], is that greater interactivity causes greater learning. Previous studies, however, have suggested that when the instructional content is strictly equivalent between conditions, highly interactive tutoring methods (such as human tutoring) are no more effective than moderately interactive methods (such as step-based NL tutoring systems) [VanLehn et al., 2007a, VanLehn, 2009]. Results such as found in this thesis suggest that more interactivity does not necessarily guarantee more learning.

Previous researchers in the cognitive and learning sciences have identified a number dialogue, domain, and student features that may be relevant when making pedagogical decisions. For the present studies, six categories of features were defined based upon prior research. These categories are discussed in Chapter 6. The final sets of tutorial tactics (used in the NormGain- and InvNormGain-Cordillera) employed more of the problem-solving context features than any other categories. The features in this category occurred 58 out of 117 times. *ConceptDifficulty*, which describes the difficulty level of the present tutorial decision, was employed by 12 of the 34 tutoring policies used in Study 3 and was the most common

single feature.

9.2 CONTRIBUTIONS TO AI&ED, ITS & EDM

The development of Intelligent Tutoring Systems (ITSs) is typically viewed as a standard system design problem. Every aspect of the system, including any pedagogical tutorial tactics employed, are specified a-priori by system developers under the guidance of or in response to domain experts. The tutoring behavior is similarly viewed as a sequential decision process where, at each discrete step, the tutor is responsible for selecting the next action to take. As with other features of the system these decisions are guided by a-priori policies. Tutoring systems face a number of limitations to their deployment, notably, the cost and pace of development make it difficult to quickly deploy systems or to adapt them to changing needs.

Recently, researchers have turned to the development of authoring tools to address these problems [Aleven et al., 2006, Ainsworth and Fleming, 2005, Murray et al., 2003]. These tools provide content support and tool libraries that permit domain experts to design and deploy systems for their needs without extensive software development. Most authoring tools, however, focus on supporting the adaptation of new content and, as such, use static predefined pedagogical strategies that allow only small amounts of tailoring by the author, such as specifying how many hints the system will provide. Doing so means assuming that all students in all domains covered by the tutor learn best using the same strategy and that domain experts know a-priori how best to teach. However, both cognitive and learning scientists have suggested that tutors should adapt to student's needs and, as discussed above, suggested that domain experts do not always track a student's knowledge level accurately [Chi et al., 2004] or make a habit of adapting their tutorial decisions to the present context [Putnam, 1987].

This dissertation described and evaluated a general methodology suitable for ITS designers to derive effective pedagogical tutorial tactics from pre-existing interactivity data rather than, as is presently common, implementing a-priori pedagogical theories drawn from experts. This approach does not require ITS designer to have an a-priori belief about how

he/she thinks the tutor should teach. Instead, this is a data-driven architecture. Having an ITS that makes effective decisions more automatically is important for its acceptance in a broader community. Systems that may be adapted to an instructor’s preferred pedagogy may be more easily acceptable to pedagogical experts. And, by applying RL to induce tutorial tactics, systems may be developed that are neither equipped with nor restricted to potentially faulty a-priori tutorial tactics.

The results herein showed that it is possible and practical to develop a strategy-neutral tutoring system (Random-Cordillera) and to use it to explore the relevant features of a tutoring space. The results also show that it is possible to use existing pedagogical knowledge (e.g. lists of relevant features) to structure the RL process and to use the results of that process to assess the domain theories.

9.3 FUTURE WORK

Given the diverse contributions of this dissertation, the work here is a starting point for additional analyses focusing on what makes this approach successful and whether it can be transferred to other domains, other machine learning methods, and other applications. The first branch of research will focus on the tutorial corpora. In the work for this dissertation policies were induced from three training corpora: Exploratory, DichGain and Combined. The final tactics used were drawn from all three corpora as discussed in Chapter 6. Among them, 11 out of 17 NormGain tutorial tactics were from the Exploratory Corpus while 11 out of 17 InvNormGain ones were from the DichGain Corpus. Given the effectiveness of NormGain relative to the Exploratory and DichGain and the null effect among the InvNormGain, Exploratory and DichGain groups, the initial analysis suggests that the Exploratory Corpus was a more effective source of tutorial tactics than the DichGain or Combined Corpora. However, it should be noted that the DichGain tutorial tactics were also derived from the Exploratory Corpus. Therefore, it can be argued that whether a training corpus is effective or not also depends on many other factors such as feature selection approach, the feature used to represent the state and so on. However this selection was based upon the estimated

success of each set of tutorial tactics and not on a full study comparison. This has important implications for the potential success of these policy induction methods with existing datasets. In future work more direct comparisons between the tutorial tactics derived from random corpora and those from preexisting corpora are planned.

A second line of proposed research will be to focus on the divergence between KC-general and KC-specific tutorial tactics. Here it was assumed that KC-specific tutorial tactics would be more effective to improve learning than KC-general ones. However, annotating everything with corresponding KCs is very time-consuming. A question exists as to whether a KC-general policy would be just as effective. For example, in Study 3, the NormGain-Cordillera and InvNormGain-Cordillera followed KC-specific tutorial tactics on eight main KCs and KC-general ones for the remainder of the 23 KCs. Later results on KC-based learning gains indeed showed that the NormGain students also learned significantly more than the InvNormGain ones on some of these 23 KCs. So it suggests that KC-general policies may be as or more effective than KC-specific tactics. Further comparisons are needed to analyze this.

In Studies 2 and 3 the state representation had been composed from features that had been suggested in the cognitive and learning science literatures and which could be both automatically computed and unambiguously evaluated. The use of manual features is impractical in the current study due to the emphasis on online training. However, there are a number of potentially relevant features that do not meet these criteria such as motivation [Hume et al., 1995, Noe, 1986]. Future work may investigate manual features or features that do not allow for unambiguous classification so much as estimation. For feature selection procedures, although there was some success here, one of the best final induced policies was from Random feature selection. So further explorations on feature selections are needed.

In addition to examining alternate feature selection policies the feature selection criteria should be studied. In this study alternate candidate features were selected based upon their Expected Cumulative Reward (ECR). In theory ECR values, as estimated from the underlying models, are a reliable indicator of future performance. However this contention has not been empirically verified in the present domain. The tactics employed in both the DichGain- and the InvNormGain-Cordillera systems, for example, have comparatively high

ECR values. However they did not seem to either outperform or fall below the Exploratory group's performance. Future work may explore the quality of ECR and investigate ways to evaluate new instructional policies. While the gold standard is to collect a new dataset by executing the new policy, it may be possible to demonstrate that cheaper evaluation methods provide sufficient accuracy to guide development. For example, one alternate policy estimation method is to test them on simulated students. Alternatively, cross-validation on ECR, e.g. train on $n-1$ students and test predicted vs. actual reward for the other one, can be run.

Additionally, the induction model used in this research was based exclusively on Markov Decision Processes (MDPs). While this framework seems effective, there are alternative methods such as Partially Observable Markov Decision Process (POMDP) [[Hauskrecht, 1997](#), [Aström, 1965](#)], which may be more suitable. POMDPs allow for realistic modeling of the student's knowledge levels, the student's intentions, and other hidden state components by incorporating them into the state space. POMDPs explicitly represent two sources of uncertainty: non-determinism in the control process and partial observability of the students' knowledge levels. In the former case, outcomes of the tutorial actions or student's knowledge level are not deterministic. In the latter, the underlying student's knowledge levels are observed indirectly via incomplete or imperfect observations.

Another future study would be to investigate why the NormGain students learned more effectively than the other three groups. One approach would be to investigate whether, by using the number of tutors' decisions on each KC, combined with the KC's I-ratio and J-ratio, its cognitive difficulty, and student's pre-test scores, one is able to predict the student's learning gains. An additional study would be to investigate when and how different RL techniques can provide increased leverage when compared to other learning techniques or be combined with other techniques (such as learning decomposition [[Beck and Mostow, 2008](#)]).

Finally, comparing these induced policies with existing learning and cognitive theories or with real human tutors in a more formal way such as evaluating and comparing them in-vivo studies could be a fruitful area of study. Previous research in non-tutoring dialogue systems has shown that the induced policy can sometimes beat human generated policies [[Lemon et al., 2006](#)]. Given the computing power of the computers, it is expected that the

induced tutorial tactics might be more effective than human tutors’.

APPENDIX A

KNOWLEDGE COMPONENTS

In this appendix I list all 32 domain KC's grouped by topic.

Table A1: Individual (not net) forces

ID	Knowledge Component (KC)
KC1	If an object is near a planet, and then the planet exerts a gravitational force on the object. The force is straight down.
KC2	The magnitude of the gravitational force is $m * g$, where m is the mass of the object and g is the gravitational acceleration of the planet.
KC3	If a spring pushes on an object, then it exerts a force on the object. The force's direction is from the spring toward the object.
KC4	If a surface pushes on an object, it exerts a normal force on the object. The force is perpendicular to the surface and away from it.
KC5	If an object contacts a surface while moving along it, and the surface is not known to be frictionless, then the surface exerts a dynamic friction force on the object. The force is parallel to the surface and in the opposite direction of the object's motion relative to the surface.
KC6	If object A pushes or pulls on object B, and the force is not one of the types listed above, then there is an applied force on B due to A. If it is a push, then it is directed from A toward B. If it is a pull, then it is directed from B toward A.

Table A2: Choosing a system for COE

ID	Knowledge Component (KC)
KC7	When an object A exerts a force on an object in the system, and the force is not a spring force nor a gravitational force, then object A must be outside the system.
KC8	When an object in the system is moving, and a spring or gravitational force acts on it and is perpendicular to its motion, then put the object exerting the force outside the system.
KC9	When an object in the system is moving, and a spring or gravitational force acting on it is not perpendicular to the moving object, then put the object exerting the force inside the system.

Table A3: Individual (not net) work

ID	Knowledge Component (KC)
KC10	If an object moves while a force interacts with it, then the force does work on the object, although the work may be zero.
KC11	If an object does not move while a force interacts with it, then the force does no work on it. (Not in textbook, but easily inferred.)
KC12	If a force does work on an moving object, and the force on an object is always perpendicular to its motion, then the work is zero. (This was generalized from the textbook version for the sake of the roller coaster training problem, 81).
KC13	If a constant force does work on an object which is moving in a straight line, and the object moves anti-parallel to the force, then the work is $-F * d$, where F is the magnitude of the force and d is the displacement.
KC14	If a constant force does work on an object which is moving in a straight line, and the object moves parallel to the force, then the work is $F * d$, where F is the magnitude of the force and d is the displacement.
KC15	The unit for work is the Joule (J) – used only when answer is work in joules.

Table A4: Net work

ID	Knowledge Component (KC)
KC16	Suppose two objects interact via a force. If one object is inside the system and one is outside, then the force is an external force. If they are both inside the system, then the force is an internal force.
KC17	If there is only one external force during a time interval, then the work done by it is also the net work done on the system during the time interval.
KC18	If there are multiple external forces acting on objects in a system, then the net work on the system is the sum of the work done by these forces on objects in the system.

Table A5: Individual (not net) mechanical energies

ID	Knowledge Component (KC)
KC19	If an object is not moving, then it has no kinetic energy.
KC20	If an object is moving, then its kinetic energy at a time is $0.5 * m * v^2$, where m is the object's mass and v is the magnitude of the object's instantaneous velocity.
KC21	If an object and a planet are in a system (or equivalently, the gravitational force of the earth on the object is an internal force), then their gravitational potential energy is $m * g * h$, where m is the mass of the object, g is the gravitational acceleration of the planet, and h is the object's height above a zero point. The zero point is arbitrary, but is often chosen to be the planet's surface.
KC22	If an object and a spring are in a system (or equivalently, the force the spring exerts on the object is an internal force), then their spring potential energy is $0.5 * k * d^2$, where k is the spring's spring constant and d is the displacement of the object relative to the equilibrium position of the spring.
KC23	The unit for energy is the Joule (J) – used only when answer is in Joules.

Table A6: COE, TME and isolated/non-isolated

ID	Knowledge Component (KC)
KC24	<p>The total mechanical energy of a system is the sum of the kinetic energies of each object in the system plus the sum of the potential energies (gravitational and spring) of each pair of objects in the system that have potential energies. Typically there is only one moving object and at most one pair for each type of potential energy, so $TME = KE + GPE + SPE$.</p>
KC25	<p>Given a system composed of one or more objects, it is isolated if there are no physical interactions between an object inside the system and an object outside the system where forces are one kind of physical interaction.</p>
KC26	<p>Otherwise, the system is non-isolated.</p>
KC27	<p>If a system isolated during a time interval, then its total mechanical energy at the beginning of the time interval equals its total mechanical energy at the end.</p>
KC28	<p>If a system is not isolated during a time interval, and forces are the only physical interactions between objects inside the system and objects outside the system, then its total mechanical energy at the beginning of the time interval equals its total mechanical energy at the end plus the net work done on the system during the time interval.</p>

Table A7: Kinematics

ID	Knowledge Component (KC)
KC29	When an object is in projectile motion (only gravity acts on it), at the apex of its trajectory, the vertical component of its velocity is zero and the magnitude of its velocity is minimal. Also, it slows down as it climbs and speeds up as it falls.
KC30	When an object slows to a stop and reverses direction, then its velocity is momentarily zero.
KC31	The unit for velocity is $\frac{m}{s}$. Used only when the answer is a numerical velocity.
KC32	The value of g is $9.8 \frac{m}{s^2}$.

APPENDIX B

GRADING

As described earlier, in studies 1 and 2, all tests were graded by a single experienced grader whom I will refer to as the *original grader*. I will likewise refer to the grades she assigned as the *original grades*. The original grades were used to calculate the KC-based NLGs and cumulative KC-based NLGs for the Exploratory and Dichotic Gain (DichGain) Corpora later used to derive the tutorial tactics for studies 2 and 3. The results I reported in chapters 4 and 5 were all based on the original scores. For study 3, however, it was necessary to engage a new grader, in this case me. I have expertise in the domain and have graded exams in this domain for a previous experiment. Since the original scores were used as the basis for our reward function it is important for us to use an equivalent rubric when grading the final exams. In order to ensure that these rubrics can be adequately aligned, I re-graded all of the original exams from studies 1 and 2 so as to conduct a full grader agreement study.

B.1 GRADING PROCEDURE

Across three studies in this dissertation, we have total 158 participants and the pre- and post-tests were identical. In order to aid grading, I first developed a simple interface to input a correct answer for each test question. As in the original grading process, each test answer was assigned two grades: overall and KC-based grading. The overall grade was a

score in the range $[0, 1]$ describing the correctness of the answer as a whole. The KC-based grading assigned a score in the range $[0, 1]$ indicating whether the student applied each relevant KC in their answer and, if so, how correct their application of the KC was. Over the course of the study we collected $64 \times 2 = 128$ students' answers for *Q22* in study 1 and $(37 + 29 + 28) \times 2 = 188$ answers for *Q22** in studies 2 and 3. For the remaining test questions we obtained a total of $158 \times 2 = 316$ answers.

Two grader interfaces were developed for overall grading and KC-based grading respectively. During each grading process, all identifying information, such as the students' ID, group ID, the test (whether pre- or post-test), were concealed from the new grader and the original grades and the original graders' correct answers for each question were not presented.

The overall grading interface displays the test question statement, the correct answer input by the new grader, a student's solution and allows for the entry of an overall score for the solution. All answers for a given question were graded serially before the grader moved on to the next test question.

While overall grading did not borrow anything from the original grading, KC-based grading borrowed the list of relevant KCs for each test question from the original grading. As part of the initial grading process, the original grader identified the list of KCs that were relevant to each test question. During the new grading process, I made use of that list rather than re-mapping each test question to 31 KCs. As I stated previously, the lists identify a total of 168 KC occurrences distributed over the 33 test questions. The KC-based grading interface displays the students' answers next to the correct answer and the list of relevant KCs. The grader then entered a score for each KC separately. Again the grades were assigned serially for all student answers before moving on to the next question.

In order to avoid errors, the test answers in study 3 were graded twice by the same procedure with the answers being checked for agreement. The other test questions were graded twice, once by each grader with the original graders' grades having already been checked for errors. In order to adequately assess the new grades' reliability, we investigated the inter-grader agreement between the original grades and the new grades across the Exploratory and DichGain groups.

B.2 INTER-GRADER AGREEMENT

Both graders agreed on the correct answers for all test questions. For the present study, inter-grader agreement was calculated using Pearson's product moment correlation which yields a correlation statistic in the range $[-1, 1]$. This comparison was made both for the overall and KC-based scores across the pre- and post-tests taken from studies 1 and 2. The level of agreement on these grades was quite high on both the pre-test scores, $r = .988, p = 0.000$, and post-test scores $r = .988, p = 0.000$ across the student groups. On the Exploratory group alone, the correlation statistics were ($r = .987, p = 0.000$) on the pre-test and ($r = .977, p = 0.000$) on the post-test. For the DichGain group exam scores, the agreement was ($r = .996, p = 0.000$) on the pre-test and ($r = .993, p = 0.000$) on the post-test.

Similar results were found for the cumulative KC-based scores across the two groups in the two studies: $r = .969, p = 0.000$ (pre-test) and ($r = .977, p = 0.000$) (post-test). For the Exploratory group within study 1 the agreement was similarly high: ($r = .963, p = 0.000$) on the pre-test and ($r = .969, p = 0.000$) on the post-test. For the DichGain group in study 2, the results are ($r = .983, p = 0.000$) on the pre-test and ($r = .984, p = 0.000$) on the post-test. Similarly high correlations were found on a per-KC basis.

The high level of correlation supports the conclusion that the grades assigned by the two graders are equivalent and may thus be reliably compared. This conclusion is bolstered by the fact that our previous findings for studies 1 and 2, reported in chapters 4 and 5, still hold with the new grades.

APPENDIX C

BACKGROUND SURVEY

C.0.1 Instructions:

1. If you do not know or remember an item, enter your best guess.
2. For any item is not applicable to you, please enter N/A.
3. If you have any questions, please feel free to ask the experimenter.

C.0.2 Questions:

Which of the following courses did you complete in **high school**? To the best of your recollection, please enter the grade you received for the **highest level** of each course you completed. (For example, if you took both Algebra I and Algebra II, please indicate your grade for Algebra II.)

Did you take an **Advanced Placement(AP)** high school course in either of the following subjects? Please answer “Yes” or “No” for each course.

I have read and signed the consent form: ___

Table C1: Questions

Date: 2009-11-02 13:15:48
 User ID: 777777
 Age: -----
 Gender: ___ Male ___ Female
 Native language: ___ English ___ Other
 High School GPA: ___ out of a maximum value of: 4.0
 SAT scores: Math: ___ Reading:___ Writing: ___
(If you took the old SAT, please enter your Verbal score as Reading)
 Name of college: -----
 Year in college: ___ *(1st = freshman, 2nd = sophomore, etc.)*
 College major(s): -----
 College GPA: ___ out of a maximum value of: ___

Table C2: High School

High-school Course	Grade	Year(e.g 1999, 2006)
Algebra		
Trigonometry		
Calculus		
Physics		

Table C3: Advanced Placement

Course	Yes	No
Calculus:	---	---
Physics:	---	---

Table C4: College-level Math

Please name any **college-level math** courses you have taken, along with grades earned and years taken:

1 _____

2 _____

3 _____

4 _____

5 _____

6 _____

Table C5: College-level physics

Please name any **college-level physics** courses you have taken, along with grades earned and years taken:

1 _____

2 _____

3 _____

4 _____

5 _____

6 _____

APPENDIX D

TEXTBOOK

D.0.2.1 Introduction

This brief textbook was designed to provide you with sufficient background knowledge of physics to be able to learn something from this study. It assumes no prior college-level physics knowledge. The next major section will introduce you to various physical quantities covered in this study, and the last major section will cover topics specific to work and energy, the topic domain for this study.

A Note About Notation: For consistency's sake, this textbook shows equations and numbers that use superscripts and subscripts the same way you would type them in this study. For example, to enter an equation with exponents, you would indicate the superscripted exponents using the carat character \wedge (*shift-6*):

$$a^2 + b^2 = c^2 \rightarrow a^2 + b^2 = c^2$$

Similarly, to enter an equation with subscripts, you would indicate the subscripts using adjacent [lower-case] letters:

$$v_f = v_i + a*t \rightarrow v_f = v_i + a*t$$

To enter an expression involving a square root, use **sqrt**:

$$c = \sqrt{a^2 + b^2} \rightarrow c = \text{sqrt}(a^2 + b^2)$$

$$v_0 = \sqrt{[KE_0 / (\frac{1}{2}*m)]} \rightarrow v_0 = \text{sqrt}[KE_0 / (0.5*m)]$$



D.0.2.2 Physical Quantities

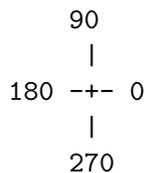
In this study you will learn about and use different physical quantities. Quantities are measurable physical features or properties, such as acceleration, length, or time, that can be expressed as variables in equations. There are two different types of quantities in physics: **scalars** and **vectors**.

A *scalar quantity* has only magnitude and is completely specified by a single numeric value with units. Some examples are length (the field is 120 yds long and 50 yds wide) and time (1 minute equals 60 s). Other examples are volume (a bottle has a volume of 1.5 liters) and frequency (electrical current has a frequency of 60 Hz). Scalar quantities of the same type can be added together using ordinary arithmetic (e.g., $60\text{ s} + 60\text{ s} = 120\text{ s} = 2\text{ min}$).

A *vector quantity* has both magnitude and direction. Some examples are acceleration (32 ft/s^2 downwards) and velocity (60 mph due east). You must be careful to account for direction when adding vector quantities. For example, when you add the velocity vectors 60 mph due east and 45 mph due west (i.e., in the opposite direction), the resulting vector is 15 mph due east not $60+45=105$ mph.

In this textbook, we will represent vector variables in bold face (**F**, **v**, **a**) and scalar quantities in standard font (m, k). Variable names will usually begin with a letter denoting its type, followed by letters or numbers distinguishing it from other variables of the same type. For example, an object's velocity at time point T1 may be represented by the variable **v1**, while the magnitude of its velocity at time T2 is v2. The mass of object A may be represented as ma, and that of object B as mb.

Whenever you are asked to enter a numerical directional orientation for a vector in this study, you should enter it in degrees relative to the right horizontal (counter-clockwise, starting from 0), according to the following convention:



The next pages in this section cover some basic physical quantities that you will encounter throughout the study.   

D.0.2.3 Mass

Mass is a scalar measure of the amount of matter that makes up an object. Its standard unit of measure is the kilogram (kg) and is often represented by a variable name beginning with m .

Mass is the property of an object that causes it to have weight in a gravitational field. An object of mass 1 kg will have 10 times as much mass as an object of mass 100 g, and it will weigh 10 times as much. However, mass is not the same as weight, which is gravitational force acting on an object. For example, an astronaut of mass 90 kg will have a different weight on the Earth than he does on the moon. We will return to this distinction later.



D.0.2.4 Displacement

Displacement is a vector quantity representing the change in position of an object. Its standard unit of measure is the meter (m) and is often represented by a variable name beginning with **d**.

It is not the same as the scalar measure of distance traveled by a moving object. For example, a swimmer who swims in a straight line from one end of a 100-m pool to the other end and back swims a total distance of 200 m, but her displacement during that same time interval is 0 m because she returns to her original position.

If a football player runs straight down the sideline from one end zone to the other (see red arrow at the bottom of the figure below), he runs a total distance of 100 yds, and his displacement is 100 yds in the direction of the far end zone. However, if he starts and ends at the same two points on the sideline as before, but runs in a zig-zag or haphazard pattern all over the field instead of a straight line (see yellow arrow below), his displacement is still 100 yds downfield but he will have run a distance greater than 100 yds.



Thus, the red arrow at the bottom of the figure (which represents the first running path) also represents the player's displacement after both running paths.

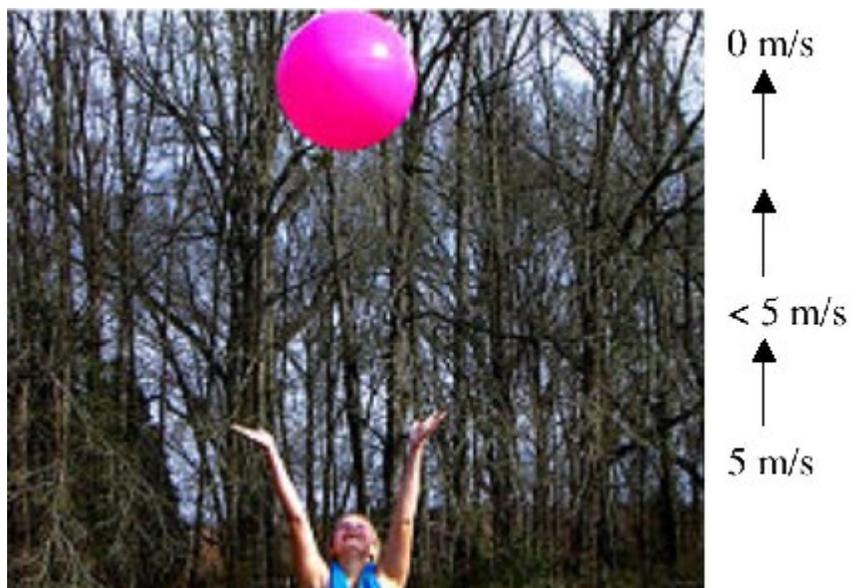


D.0.2.5 Velocity

Velocity is a vector quantity representing the rate of change in position of an object per unit time. Its standard unit of measure is meters per second (m/s) and is often represented by a variable name beginning with **v**.

The *average* velocity of a moving object is its displacement during a time interval divided by the duration of the time interval. However, velocity in this study will refer to the *instantaneous* velocity of an object, which is its displacement per unit time at a given point in time (or, over an infinitesimal time interval).

For example, let's say you throw a ball straight up in the air with an initial velocity of 5 m/s upwards. A short time later, its velocity will be less than 5 m/s upwards because gravity will slow it down. At the point in time where the ball ceases to move upwards but before it begins to fall back down, it has a velocity of 0 m/s. It will eventually land back in your hands with a downward velocity of some magnitude.



D.0.2.6 Acceleration

Acceleration is a vector quantity representing the rate of change in velocity of an object per unit time. Its standard unit of measure is meters per second per second, or meters per second squared (m/s^2) and is often represented by a variable name beginning with **a**.

The *average* acceleration of a moving object is its change in velocity during a time interval divided by the duration of the time interval. However, acceleration in this study will refer to the *instantaneous* acceleration of an object, which is its change in velocity per unit time at a given point in time (or, over an infinitesimal time interval).

As with other vector variables, be careful to account for direction when dealing with acceleration. For example, suppose the driver of a car moving east at 60 mph applies the brakes. The car's velocity vector (with magnitude v below) will continue to be eastward until the car stops, but the braking acceleration vector (with magnitude a below) will be to the *west*, in the direction opposite the car's displacement.



Although some vectors change in magnitude or direction as an object moves (such as the velocity on this and the previous page), most of the accelerations you will encounter in this study will be *constant*.

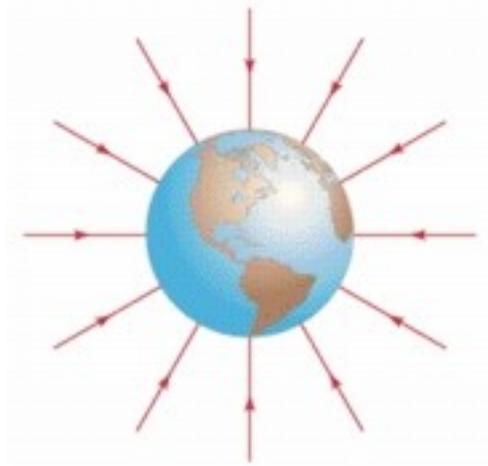


D.0.2.7 Gravitational Acceleration

One such constant acceleration used in this study is the acceleration due to gravity of a moving object near the surface of the Earth. It can be shown experimentally that all objects near the Earth's surface have exactly the same downward acceleration, whenever the effects of air resistance can be eliminated or otherwise ignored.

Unlike many other accelerations, this special constant is *not* represented by a variable name beginning with **a**. Its magnitude is represented by the scalar variable g , and its direction is always straight down (toward the center of the Earth):

$$g = 9.8 \text{ m/s}^2$$



D.0.2.8 Force

A **force** is a push or a pull exerted on an object.

It is a vector quantity and its standard unit of measure is newtons (N), which is equivalent to kilogram meters per second squared ($\text{kg}\cdot\text{m}/\text{s}^2$). It is often represented by a variable name beginning with **F**.

If a force is exerted by something in direct contact with the object, it is a *contact* force. For example, if you pull a spring-loaded pinball plunger back toward you and then let it go, it exerts a contact force on the pinball touching it, in the direction away from you. If you tie a rope around a box and start pulling it across the floor, the rope exerts a tension force on the box in the direction of your motion.

Other forces on an object that do not result from direct contact are called *field* forces. A common example is gravitation. All objects near the Earth's surface are subject to a downward force due to the Earth's gravitational field.



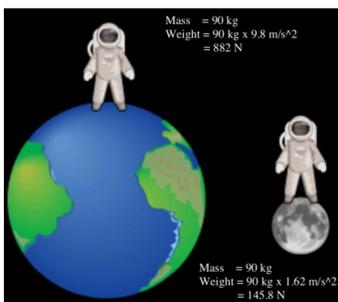
D.0.2.9 Weight Force

As discussed earlier, weight is different from mass. Weight is a force exerted on an object within a gravitational field. It is, therefore, a field force. As with all forces, the weight force is a vector quantity and its standard unit of measure is newtons(N). It is often represented by the variable **F**.

The magnitude of the weight force on an object (**F**) equals the mass of the object multiplied by the acceleration due to gravity (which is *g* if the object is near the Earth's surface). As with gravitational acceleration, the direction of the weight force vector is always straight down.

$$F = m \cdot g$$

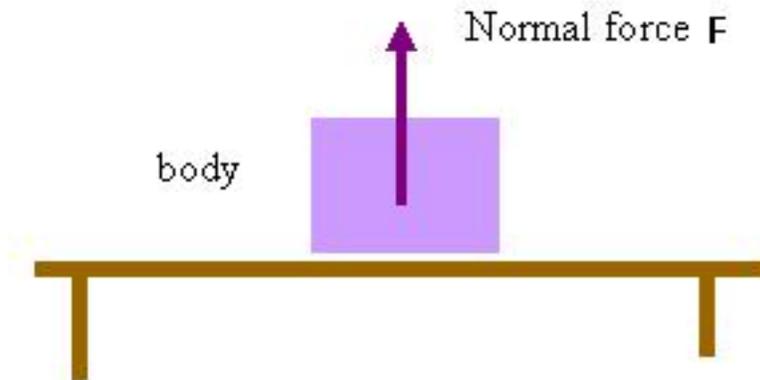
As discussed earlier, a 90-kg astronaut will have a different weight on the Earth than he does on the Moon. On the Earth he will weigh $(90 \text{ kg}) \cdot (9.8 \text{ m/s}^2) = 882 \text{ N}$, but on the Moon he will weigh only about 1/6 as much (145.8 N), because the acceleration due to gravity near the surface of the Moon is only 1.62 m/s^2 .



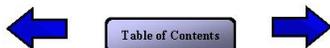
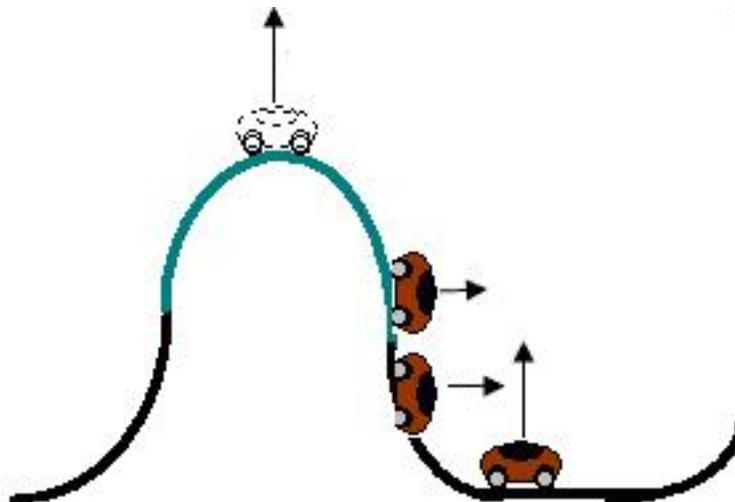
D.0.2.10 Normal Force

When an object is pressed against a surface, the object experiences an opposing contact force that is perpendicular to the surface. This force is called a *normal* force, because “normal” is a mathematical term meaning perpendicular. As with all forces, a normal force is a vector quantity and its standard unit of measure is newtons (N). It is often represented by the variable \mathbf{F} .

Suppose a rectangular object is resting on a table in your kitchen. You already know the object experiences a downward weight force, proportional to its mass. It also experiences an upward normal force from the table, which in this case is equal in magnitude and opposite in direction to the weight force on the object.

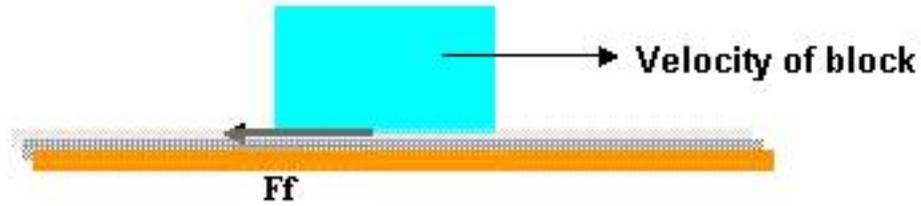


The normal force from a surface is always perpendicular to the point of contact. Therefore, if an object is moving along a curved surface (e.g., a rollercoaster car on a curved track), the direction of the normal force will change as the object moves along the curved path:



D.0.2.11 Friction Force

For simplicity's sake, many of the problem scenarios you encounter in this study will involve frictionless surfaces, or will state that friction is negligible. However, some problems will mention frictional force acting on a moving object, usually an object sliding across a surface. As with all forces, a friction force is a vector quantity and its standard unit of measure is newtons (N). A friction force (\mathbf{F}) is a force along a surface that opposes the sliding of an adjacent object across the surface. The friction force is parallel to the surface and opposite in direction to the object's motion.



D.0.2.12 Introduction to energy Energy can take many forms, but when properly defined and measured, it turns out that the total energy of an isolated system does not change over time. This law of nature is called *conservation of energy*. Notice that it uses the terms “energy” and “isolated system,” so we will need to define those carefully.

For “isolated system”, we will first define a “system” to be a set of one or more objects and then define “isolated system” as one where none of the objects inside the system interact with objects outside the system. Because some systems are not isolated, we will also show how to analyze a certain kind of non-isolated system that is particularly common.

For “energy,” we will define two forms of energy quite carefully and lump all other forms of energy into a third, “catch-all” category. The three categories of energy are

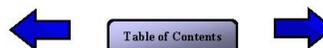
- The *kinetic energy* of a moving object. We will define this carefully.
- The *potential energy* of a pair of interacting objects. This is energy that can be easily converted into kinetic energy, which is why it is called “potential” energy. We will define this carefully.
- The *internal energy* of an object. This category covers many kinds of energy, such as temperature, chemical energy, or biological energy, that are not so easily converted to kinetic energy. We will not define this category carefully because it encompasses so many forms of energy.

Kinetic and potential energies are mechanical energies, while internal energy is not a mechanical energy. From the general conservation of energy law we stated above, it follows that, assuming the internal energy of the system does not change, then:

The total mechanical energy of an isolated system does not change.

This is a simple version of Conservation of Mechanical Energy. There is also a more general version that works for certain kinds of non-isolated systems. We will first introduce the simple version, then build up to the final, more complex version.

Because we want to focus on Conservation of Mechanical Energy, we will make the blanket assumption throughout this experiment that the internal energies of objects do not change. That is, we will assume that none of the objects in the examples, problems, or illustrations ever change their temperature, their chemical energy, etc.



D.0.2.13 Kinetic energy

One type of energy, called kinetic energy (KE), is associated with individual objects. It depends only on an object's mass and on the magnitude of its velocity. In particular, it is defined via the formula:

$$KE = 0.5 * m * v^2$$

where m is the mass of the object and v is the magnitude of the velocity of the object relative to a stationary point of reference, such as the Earth. Although the Earth has mass and moves through space, it can be considered a stationary point of reference for other objects near it, with zero velocity and, hence, zero kinetic energy.

The standard unit of energy is $(\text{kg} \cdot \text{m}^2) / \text{s}^2$ which is also denoted as J (Joule).

Example: a box of mass 2 kg travels in a circle with a velocity of 3 m/s. What is its kinetic energy at any point in the circle?

Solution:

$$\begin{aligned} KE &= 0.5 * m * v^2 \\ &= 0.5 * (2 \text{ kg}) * (3 \text{ m/s})^2 \\ &= 9 \text{ kg} \cdot \text{m}^2 / \text{s}^2 = 9 \text{ J} \end{aligned}$$



D.0.2.14 Potential energy Suppose you toss a golf ball straight up so that it reaches a peak of 10 meters before it turns around and falls to the ground. When it is on its way up, say at 2 meters off the ground, it has a non-zero velocity so $0.5*m*v^2$ is non-zero and thus it has a non-zero kinetic energy. The golf ball slows down as it climbs, so its kinetic energy decreases. At the very instant it reaches its peak, its velocity is zero for just that instant, so $0.5*m*v^2$ is zero and thus its kinetic energy is zero at that instant. As the golf ball falls down, it picks up speed, so it has a non-zero, increasing velocity, so $0.5*m*v^2$ is non-zero and increasing, and thus its kinetic energy is non-zero and increasing. In other words, while the golf ball is in flight, its kinetic energy decreases to zero then increases again.

Conservation of Mechanical Energy suggests that the golf ball's kinetic energy didn't just disappear as it rose, but instead was converted to another form of energy. As the golf ball fell, its kinetic energy was returned to it, according to Conservation of Mechanical Energy. So what is this "other form of energy" that borrows energy from kinetic energy and then returns it?

When a form of energy is easily converted from kinetic energy and back, that form of energy is called *potential* energy. Potential energy is associated with relative positions of *two* objects: in the example above, different positions of the ball with respect to the Earth account for the different values of the potential energy of the ball-Earth system.

We will not give a precise definition of "easily converted" in general nor of "potential energy" in general. However, we will define precisely two common forms of potential energy: gravitational potential energy and spring potential energy. 

D.0.2.15 Gravitational potential energy Gravitational potential energy exists when two objects are pulled together by a gravitational force. Although gravitational forces exist between any pair of objects, as long as they both have mass, the force is too small to bother with unless one of the objects is very massive (like a planet) and the other object is fairly close to it. In that case, gravitational potential energy is defined as:

$$\text{GPE} = m * g * h$$

where h is the height of the object above the level of origin, g is the constant gravitational acceleration near the surface of the planet (9.8 m/s^2 for Earth), and m is the object's mass.

Since we will only be interested in the *change* of GPE as an object moves with respect to the Earth, we can choose the level of origin to be anywhere, as long as it is fixed with respect to the Earth. For any object at the level of origin, its height $h = 0$, and therefore the potential energy of the object-Earth pair $\text{GPE} = 0$. GPE will be positive when the object is above the level of origin, and it will be negative when the object is below it.

The unit of energy is $\text{kg} * \text{m}^2 / \text{s}^2$, also denoted as J (Joule).

Example: Suppose a 0.1 kg ball is tossed straight up so that it peaks at 10 meters off the surface of the Earth. Suppose we define the origin to be 2 meters off the surface of the Earth. Thus, $h = (10 - 2) = 8 \text{ m}$ when the ball is at its peak, so the gravitational potential energy at its peak is:

$$\begin{aligned} \text{GPE} &= m * g * h \\ &= (0.1 \text{ kg}) * (9.8 \text{ m/s}^2) * (8 \text{ m}) \\ &= 7.84 \text{ kg} * \text{m}^2 / \text{s}^2 = 7.84 \text{ J} \end{aligned}$$

If the ball reaches the ground, the GPE of the ball-Earth pair at that point is:

$$\begin{aligned} \text{GPE} &= m * g * h \\ &= (0.1 \text{ kg}) * (9.8 \text{ m/s}^2) * (-2 \text{ m}) \\ &= -1.96 \text{ J} \end{aligned}$$



D.0.2.16 Spring potential energy A spring potential energy exists when one end of a spring is anchored so that it does not move, and the other end exerts a force on an object (e.g., see the *Compressed* spring in Figure 2). The value of the spring potential energy of an object-spring pair is defined as:

$$SPE = 0.5 * k * d^2$$

where k is a spring constant, and d is the extension (or compression) of the spring from its equilibrium position. The spring constant k is a positive number that measures the stiffness of a spring. The stiffer the spring, the larger the spring constant. The unit for a spring constant is N/m (which is equivalent to kg/s^2). The equilibrium position of a spring is the position of the unanchored end when the spring is not being compressed or extended. The variable d measures how far the spring is compressed or extended, and it is a positive scalar number in either case.

For the sake of simplicity, assume that all springs you encounter in this experiment are *ideal* springs that have no mass of their own and do not degrade or deform, over time or due to extended compression or extension.

The unit of energy is $kg * (m^2)/(s^2)$, also denoted as J (Joule).

Example: Figure 1 shows a spring attached to a wall at one end and a mass m at the other end. Suppose its spring constant is 0.2 N/m. It is at its equilibrium position, so the spring potential energy of the spring-mass pair is:

$$SPE = 0.5 * k * d^2 = 0.5 * (0.2N/m) * (0m)^2 = 0J$$

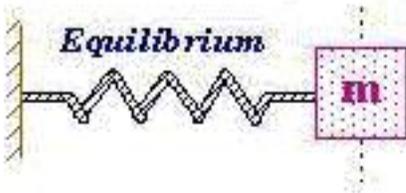


Figure 1

Figure 2 shows the spring compressed by 0.3 m. The spring potential energy of the spring-mass pair is thus:

$$SPE = 0.5 * k * d^2 = 0.5 * (0.2N/m) * (0.3m)^2 = 0.009kg * (m^2)/(s^2) = 0.009J$$

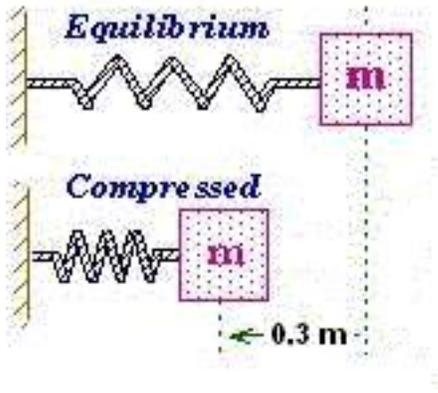


Figure 2

Figure 3 shows the spring stretched by 0.6 m. The spring potential energy of the spring-mass pair is thus:

$$SPE = 0.5 * k * d^2 = 0.5 * (0.2N/m) * (0.6m)^2 = 0.036kg * m^2/s^2 = 0.036J$$

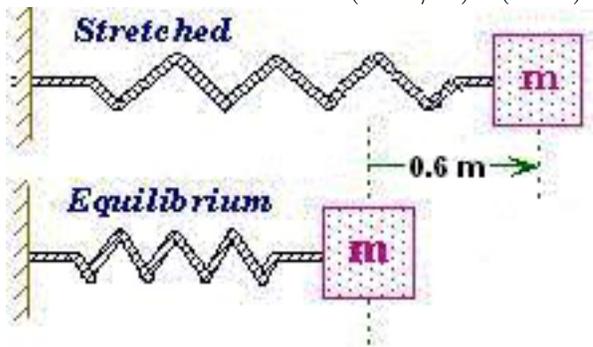


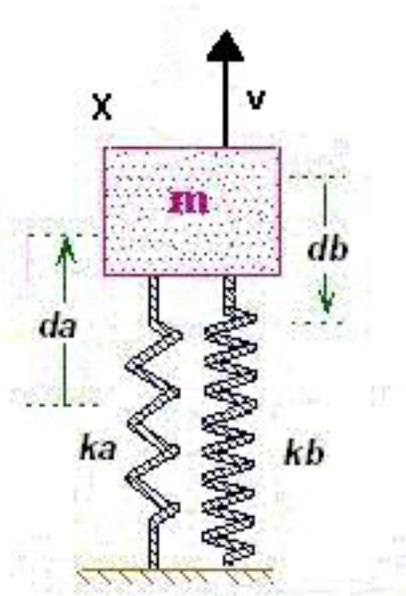
Figure 3



D.0.2.17 Total mechanical energy A *system*, as mentioned before, is just a set of one or more objects. Intuitively, the total mechanical energy of a system should be just the sum of the energies of all the objects in the system. However, potential energies are defined only for *pairs* of objects. We include a potential energy in the sum only if *both* objects are included in the system. Thus:

The total mechanical energy of a system is the sum of the kinetic energies of each object in the system plus the sum of the potential energies of each pair of objects in the system that have potential energies.

For example, consider an object X that is being pushed by two springs, Sa and Sb. The two springs are attached to a platform that is floating in deep space (i.e., nowhere near a planet). Let us consider all the possible contributions to total mechanical energy:



- Suppose that X is moving with velocity v and has mass m , so it has kinetic energy $\frac{1}{2} * m * v^2$.
- Suppose that springs Sa and Sb have no mass, so the kinetic energy of each spring is 0.
- Suppose that the spring constant of Sa is k_a and it is compressed by d_a . Thus, the potential energy of the pair X and Sa is $\frac{1}{2} * k_a * d_a^2$.
- Suppose that the spring constant of Sb is k_b and it is compressed by d_b . Thus, the potential energy of the pair X and Sb is $\frac{1}{2} * k_b * d_b^2$.

Now we can demonstrate that the total mechanical energy (i.e., the sum of each object's kinetic energy plus any potential energies of object pairs) depends on which set of objects comprise the system:

- The total mechanical energy of {X, Sa, Sb} is $\frac{1}{2} * m * v^2 + 0 + 0 + \frac{1}{2} * ka * da^2 + \frac{1}{2} * kb * db^2$
- The total mechanical energy of {X, Sa} is $\frac{1}{2} * m * v^2 + 0 + \frac{1}{2} * ka * da^2$
- The total mechanical energy of {X, Sb} is $\frac{1}{2} * m * v^2 + 0 + \frac{1}{2} * kb * db^2$
- The total mechanical energy of {X} is $\frac{1}{2} * m * v^2$

Similarly, gravitational potential energy is defined between a pair of objects, one of which is usually a planet. For instance, if a rock of mass m is near the Earth and moving with velocity v (relative to the Earth, which is to be considered stationary), then the total mechanical energy of {rock, Earth} is $\frac{1}{2} * m * v^2 + m * g * h$, where h is the height of the rock above a reference point. The total mechanical energy of {rock} is $\frac{1}{2} * m * v^2$, because the Earth is not in the system.



D.0.2.18 Isolated and non-isolated systems A system is called *isolated* whenever all of the objects in it have no physical interaction with objects outside it. If at least one object in a system has a physical interaction with some object outside the system, then the system is called *non-isolated*.

There are many kinds of physical interactions. One kind is a force that exists between a pair of objects. For instance, if object A pushes object B, then A exerts a force on B. If A is in the system and B is not, then the system is non-isolated. In particular, if object P is a planet that exerts a non-negligible force on object O, then a system that has O but not P is non-isolated.

For example, consider an object falling toward the Earth in a vacuum. Because it is in a vacuum, nothing is touching it, so the only force on it is the gravitational force of the Earth. If we define the system to be {object, Earth}, then it is isolated because nothing is interacting with the objects in it. If we define the system to be just {object}, then the system is non-isolated, because the Earth is exerting a force on the object.

There are also physical interactions between objects that are not due to forces. For instance, if one object radiates heat and warms up another object, then there is no force between them, but there is a physical interaction.



D.0.2.19 Conservation of Mechanical Energy for isolated systems Having defined the key terms “total mechanical energy” and “isolated system”, we can now return to the simpler version of the law of Conservation of Mechanical Energy, which applies only to isolated systems:

The total mechanical energy of an isolated system does not change.

Example: Suppose at time 1, a 0.1 kg ball is launched vertically upward from the surface of the Earth with a velocity of 5 m/s. At time 2, it reaches its peak of 1.28 meters above the surface of the Earth and its velocity is momentarily zero.

To analyze this situation with Conservation of Mechanical Energy, let us consider the ball and the Earth to be a system. In order to keep the example simple, let us ignore air, wind, ambient sound, sunlight, and all the other physical interactions that actually exist, and treat this two-object system as isolated.

The law says that the total mechanical energy of {ball, Earth} at time 1 should equal the total mechanical energy of {ball, Earth} at time 2. Let’s check: At time 1,

$$ME1 = KE1 + GPE1$$

where ME1 is the total mechanical energy at time1, KE1 is the total kinetic energy at time 1, and GPE1 is the total gravitational potential energy at time 1. There is no spring potential energy because there are no springs in the example. At time 2,

$$ME2 = KE2 + GPE2$$

where these variables refer to the total mechanical energy, kinetic energy, and gravitational potential energy at time 2. Now let’s consider each type of energy. For KE1, the Earth is stationary but the ball starts with a velocity of 5 m/s. Thus, $KE1 = 0 + 0.5 * m * v^2 = 0.5 * (0.1 \text{ kg}) * (5 \text{ m/s})^2 = 1.25 \text{ J}$. For KE2, the Earth is stationary and the ball has zero velocity as well. Thus, $KE2 = 0 + 0 = 0$.

For gravitational potential energy, we need to choose an origin so we can calculate a height above it. For simplicity’s sake let us choose the surface of the Earth as the origin. Thus, for GPE1, $m * g * h1$ is 0 because the height ($h1$) is zero; therefore, $GPE1 = 0$. For $GPE2 = m * g * h2$, we have $h2 = 1.28 \text{ m}$, so $GPE2 = (0.1 \text{ kg}) * (9.8 \text{ m/s}^2) * (1.28 \text{ m}) = 1.25 \text{ J}$.

Now we can check Conservation of Mechanical Energy; that is, does $ME1 = ME2$? We simply plug in the values we have calculated:

$$\begin{aligned}ME1 &= KE1 + GPE1 \\ &= 1.25 \text{ J} + 0 \\ ME2 &= KE2 + GPE2 \\ &= 0 + 1.25 \text{ J}\end{aligned}$$

This shows that $ME1=ME2$. Essentially, the kinetic energy at time1 has been transformed into gravitational potential energy at time2.   

D.0.2.20 Internal forces vs. external forces A system is non-isolated if there exists at least one physical interaction between objects in the system and objects outside the system. We can extend the law of Conservation of Mechanical Energy to handle a special class of non-isolated system, where all the physical interactions are due to forces. In order to do that, we first need to define several new terms: *internal/external forces*, *work*, and *net work*. This page defines internal and external forces.

Recall that a force is defined only for a *pair* of objects: the object that the force acts on, and the object that is exerting the force. For instance, when the Earth exerts a gravitational force on a ball, there are two objects involved: the Earth and the ball. When a frictional force slows down a sliding block, the surface (one object) is exerting a force on the sliding block (the second object).

If there is a force involving a pair of objects, and we define the system to include only one of them, then the force is crossing the system boundary and the system is non-isolated. We call such a force an *external* force.

On the other hand, if both objects of the pair are inside the system, then we call the force an *internal* force.



D.0.2.21 Work The previous page defined internal and external forces, but we still need to define *work* in order to handle the afore mentioned special class of non-isolated system where all the physical interactions are due to forces. A completely general definition of work can be stated using vector calculus, but we will not do so here. Instead, we will define work only for a few rather common configurations.

Work is defined for a particular object and a particular force over a particular time interval. We will assume that through the time period, (1) the force is constant and has magnitude F , and (2) the object moves in a straight line and its displacement has magnitude d . Given these assumptions, the work W done on the object by the force depends on the relative direction of the force and the displacement:

- If the object is moving in the *same* direction as the force, then $W = F \cdot d$ (*positive*).
- If the object is moving in the *opposite* direction to the force, then $W = -F \cdot d$ (*negative*).
- If the object is moving *perpendicular* to the force, then $W = 0$.

Differences in how we handle the work done by internal versus external forces will be discussed on later pages (e.g., regarding different choices of systems).   

D.0.2.22 Net work We have defined work for a single object, but Conservation of Mechanical Energy applies to a system, which is a set of one or more objects. Thus, we need to define the *net work done on a system* as the sum of the work done by the external forces on objects inside the system. Notice that we only sum over the *external* forces and do not include the internal ones.

For example, consider a crane that is lifting a block against gravity. If we define the system to be {Earth, block} so that the crane is outside the system, then the force exerted by the crane on the block is external whereas the gravitational force is internal. The net work done on the system is the work done by the crane on the block. If the crane exerts a force of magnitude F and lifts the block a displacement of magnitude d , then the work done by the crane on the block (W) is $F*d$, so the net work on the system (W_{net}) is also $F*d$.

On the other hand, if we define the system to be {block}, then both the crane and the Earth are outside the system, so there are two external forces. As before, the work on the block due to the crane is $F*d$. If we use F_w for the magnitude of the weight force due to gravity, then the work done by the Earth is $-F_w*d$. (Note the minus sign, because the weight force is downward but the block's displacement is upward.) Thus, the net work done on the system is $W_{net}=F*d-F_w*d$.



D.0.2.23 Conservation of Mechanical Energy for systems whose non-isolation is due to forces Now we can extend the law of Conservation of Mechanical Energy to systems which are not isolated, but whose non-isolation is due to forces and where no other kind of physical interaction (e.g., light, sound, heat) exists. In this case, the law is

$$ME1 + W_{\text{net}} = ME2$$

where $ME1$ is the total mechanical energy at time 1, $ME2$ is the total mechanical energy at time 2, and W_{net} is the net work done on the system during that time interval.

Example: Suppose a little girl on ice skates is given a slow push by her father, who is standing still. The magnitude of the force exerted by the father is 0.2 N, and the distance traveled by the little girl while she is being pushed is 0.5 m. What kinetic energy does the little girl attain, starting from rest?

We can solve this problem by using a non-isolated system comprised of the little girl and the Earth. We will treat the ice and the father as being outside the system. First, let us classify all the forces in this situation as either internal or external, and compute the work done by the external ones.

- The father exerts a force on the girl. It is external, and so the work done is the force's magnitude (0.2 N) times the magnitude of the girl's displacement (0.5 m). Thus, the work done is 0.10J.
- The ice exerts a normal force on the girl. It is external. The force is straight up, and the girl moves horizontally, so the force and the displacement are perpendicular. Thus, the work done by the normal force on the girl is 0.
- The ice exerts a tiny friction force on the girl. Although it is external, it is so small that we can ignore the work done by it.
- The Earth exerts a gravitational force on the girl. This force is internal, because we are including the Earth in the system.

Summing up the net work done during the time interval on the {girl, Earth} system, we have $W_{\text{net}}=0.10$ J.

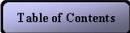
Initially, the system's total mechanical energy, $ME1$, is $KE1 + GPE1$. After the push ends, the system's energy, $ME2$, is $KE2 + GPE2$. Substituting these into conservation of

energy, we have

$$KE_1 + GPE_1 + W_{\text{net}} = KE_2 + GPE_2$$

The girl has the same height as she did initially, so $GPE_2 = GPE_1$. Thus, we can simplify the conservation of mechanical energy equation to

$$KE_1 + W_{\text{net}} = KE_2$$

We know that $KE_1 = 0$ because neither the girl nor the Earth are moving initially. We know that KE_2 is just the kinetic energy of the girl, because the Earth is still not moving. After substituting these relationships for KE_1 and KE_2 into $KE_1 + W_{\text{net}} = KE_2$, we obtain $W_{\text{net}} =$ the girl's kinetic energy, so her kinetic energy is 0.10 J.   

D.0.2.24 Different choices of system On the preceding page, we chose {girl, Earth} as the system, making the Earth's gravitational force internal to the system. Suppose we instead chose {girl} as the system. Now the Earth is outside the system, so that its gravitational force becomes an external force.

Because we have added an external force, the net work done on the {girl} could be different than the net work done on the {girl, Earth}. However, as it turns out, the gravitational force is perpendicular to the girl's displacement, so the work done by it on the girl is zero. Thus, the net work done on {girl} is the same as the net work done on {girl, Earth}.

There is a second change in the analysis. With {girl, Earth} as the system, the total mechanical energies, ME1 and ME2, include the girl/Earth gravitational potential energies, GPE1 and GPE2. With {girl} as the system, the gravitational potential energies are excluded because potential energies exist only between *pairs* of objects in a system. Thus, applying the law to {girl, Earth}, we got $KE1 + GPE1 + W_{net} = KE2 + GPE2$, whereas applying the law to {girl}, we now get $KE1 + W_{net} = KE2$. However, $GPE1 = GPE2$ with {girl, Earth}, so the solutions converge and we get exactly the same answer for both choices of system.

What happens if we choose {father, girl, Earth} as the system? This would make the father's push into an internal force (it was external). Hence, the net work on the 3-object system would be zero (it was $F \cdot d$, the work done by the father on the girl). With a little more math, we can show that the girl's kinetic energy would be zero, which is wrong! Thus, we do *not* have total freedom when choosing a system. Only some choices give correct answers.

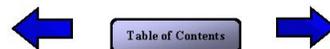
In order to know which choice of system to make, we must first discuss the differences between conservative and non-conservative forces.   

D.0.2.25 Conservative and non-conservative forces A force is called *conservative* if and only if we can define a potential energy for it. So far, we have defined potential energies for gravitational forces and for spring forces, so those are the only conservative forces introduced so far. All the other forces introduced so far (normal force, frictional force, pulls by a crane, pushes by fathers, etc.) are *non-conservative* forces.

Recall that the potential energy of a pair of objects depends only on their relative position. Thus, if one object of the pair moves around during a time interval but comes back to exactly the same position, then the potential energy of the pair has not changed during the time interval. For the non-conservative forces, there is no definition of potential energy that has this property, and that is why they are non-conservative.

For instance, suppose we tried to define a potential energy for the father's force on his little girl. If there was such a "father's force potential energy," then we would include it when summing up total mechanical energy. Now suppose at time 1 the girl is not moving. Next the father pushes her away and pulls her back a hundred times. At time 2, she is back at the original position and not moving. Now her kinetic energy at both times is zero, her gravitational potential energy has not changed, and moreover, her "father's force potential energy" has not changed either, because she is back where she started. Thus, her total mechanical energy, which includes the "father's force potential energy," has not changed. Yet we know he will have expended some biological energy (i.e., burned some calories), and thus his internal energy has changed. Thus, we have a change in internal energy without a compensating change in total mechanical energy. This violates the general law of Conservation of Energy. Thus, it is impossible to define a potential energy for the father's push that will allow the conservation laws to hold.

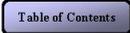
For this experiment, you can assume that the only conservative forces are gravitational and spring forces. Although a few other conservative forces are known in physics, they will not appear in any of the examples or problems used in this experiment.



D.0.2.26 Internal forces must be conservative When choosing a system, *all the internal forces must be conservative*. As demonstrated earlier, when a non-conservative force such as the father's push is included as an internal force, the analysis is incorrect.

When we gather up all the various caveats that have been discussed, a final version of the law of Conservation of Mechanical Energy can now be stated:

If a system can be chosen such that (a) all the internal forces are conservative, (b) external forces are the only physical interactions between objects inside the system and objects outside the system, and (c) no objects inside the system change their internal energy, then $ME_1 + W_{net} = ME_2$, where ME_1 is the total mechanical energy at time 1, W_{net} is the net work done on the system by external forces between times 1 and 2, and ME_2 is the total mechanical energy of the system at time 2.

When the system is isolated, then there are no external forces, so the net work is zero and so $ME_1 = ME_2$. Thus, the first version of Conservation of Mechanical Energy we introduced, which applied only to isolated systems, is clearly just a special case of this more general version.  

D.0.2.27 Summary

Below is a summary table of all the major work and energy principles covered in this textbook. Please review it now before continuing the next phase of the experiment.

PRINCIPLE	EQUATION
Kinetic energy	$KE = 0.5 * m * v^2$
Gravitational potential energy	$GPE = m * g * h$
Spring potential energy	$SPE = 0.5 * k * d^2$
Total mechanical energy	$ME = KE + GPE + SPE$
Work	$W = \mathbf{F} * \mathbf{d}$
$W = F * d$	<i>if same direction</i>
$W = -F * d$	<i>if opposite direction</i>
$W = 0$ if	<i>perpendicular</i>
Net work	$W_{net} = \Sigma W = \text{sigma}(W)$ (<i>= sum of all Ws</i>)
Conservation of mechanical energy	$ME_0 + W_{net} = ME_1$

($W_{net}=0$ if isolated system)



Table of Contents

APPENDIX E

PRE- AND POSTTEST QUESTIONS.

Please enter the answer to the following question in the space below. (Value: 2/100)

1. Enter the equation that defines the kinetic energy of an object (remember to use * for multiplication and $\hat{}$ for exponentiation, and be sure to include an = sign):

Please enter the answer to the following question in the space below. (Value: 2/100)

2. Enter the equation that defines the gravitational potential energy of an object (remember to use * for multiplication and $\hat{}$ for exponentiation, and be sure to include an = sign):

Please enter the answer to the following question in the space below. (Value: 2/100)

3. Enter the equation that defines the spring potential energy of an object (remember to use * for multiplication and $\hat{}$ for exponentiation, and be sure to include an = sign):

Please enter the answer to the following question in the space below. (Value: 2/100)

4. During a time interval, an object moves in a straight line, a constant force acts on it, and the force is in the same direction as the object's motion. Enter the equation that defines the work done on an object by the force (remember to use * for multiplication and $\hat{}$ for exponentiation, and be sure to include an = sign):

Please enter the answer to the following question in the space below. (Value: 2/100)

5. Enter the equation that defines the total mechanical energy of a system (remember to use * for multiplication and $\hat{}$ for exponentiation, and be sure to include an = sign):

Please select ALL of the possible answers to the following question from the list. (Value: 2/100)

6. Suppose an object is near Earth, moving, attached to a spring and acted on by an applied force F . What does the kinetic energy of the object depend on directly? That is, what is mentioned in the definition of kinetic energy? Check all that apply:

1. The mass of the object.
2. The magnitude of the object's acceleration.
3. The direction of the object's acceleration.
4. The magnitude of the object's velocity.
5. The direction of the object's velocity.
6. The magnitude of the object's displacement.
7. The direction of the object's displacement.
8. The magnitude of the force F .
9. The direction of the force F .
10. The spring constant of the spring.
11. The distance that the spring is compressed or extended.
12. The gravitational constant of the Earth.
13. The height of the object above the surface of the Earth.

14. The absence of any other physical interactions between the object and its environment other than the ones mentioned above (the spring, the Earth and the applied force F).

Please select ALL of the possible answers to the following question from the list. (Value: 2/100)

7. Suppose an object is near Earth, moving, attached to a spring and acted on by an applied force F . What does the gravitational potential energy of the object depend on directly? That is, what is mentioned in the definition of gravitational energy? Check all that apply:

1. The mass of the object.
2. The magnitude of the object's acceleration.
3. The direction of the object's acceleration.
4. The magnitude of the object's velocity.
5. The direction of the object's velocity
6. The magnitude of the object's displacement.
7. The direction of the object's displacement.
8. The magnitude of the force F .
9. The direction of the force F .
10. The spring constant of the spring.
11. The distance that the spring is compressed or extended.
12. The gravitational constant of the Earth.
13. The height of the object above the surface of the Earth.
14. The absence of any other physical interactions between the object and its environment other than the ones mentioned above (the spring, the Earth and the applied force F).

Please select ALL of the possible answers to the following question from the list. (Value: 2/100)

8. Suppose an object is near Earth, moving, attached to a spring and acted on by an applied force F . What does the spring potential energy of the object depend on directly? That is, what is mentioned in the definition of spring potential energy? Check all that apply:

1. The mass of the object.
2. The magnitude of the object's acceleration.
3. The direction of the object's acceleration.
4. The magnitude of the object's velocity.
5. The direction of the object's velocity.
6. The magnitude of the object's displacement.
7. The direction of the object's displacement.
8. The magnitude of the force F .
9. The direction of the force F .
10. The spring constant of the spring.
11. The distance that the spring is compressed or extended.
12. The gravitational constant of the Earth.
13. The height of the object above the surface of the Earth.
14. The absence of any other physical interactions between the object and its environment other than the ones mentioned above (the spring, the Earth and the applied force F).

9. Suppose an object is near Earth, moving, attached to a spring and acted on by an applied force F . What does the work done by F on the object depend on directly? That is, what is mentioned in the definition of work done on an object? Check all that apply:

1. The mass of the object.
2. The magnitude of the object's acceleration.
3. The direction of the object's acceleration.
4. The magnitude of the object's velocity.
5. The direction of the object's velocity.
6. The magnitude of the object's displacement.

7. The direction of the object's displacement.
8. The magnitude of the force F .
9. The direction of the force F .
10. The spring constant of the spring.
11. The distance that the spring is compressed or extended.
12. The gravitational constant of the Earth.
13. The height of the object above the surface of the Earth.
14. The absence of any other physical interactions between the object and its environment other than the ones mentioned above (the spring, the Earth and the applied force F).

Please enter the answer to the following question in the space below. (Value: 4/100)

10. A toy cart moves with a kinetic energy of 30 J. If the magnitude of its velocity is doubled, what will its kinetic energy be?

Please enter the answer to the following question in the space below. (Value: 4/100)

11. A force of 500 N is exerted on a given object. How far can the object move if 5000 J of work are available to move it?

Please enter the answer to the following question in the space below. (Value: 4/100)

12. A baseball with mass 0.145 kg is thrown straight up with an initial velocity of magnitude 20.0 m/s. What is the work done by gravity on the baseball when it reaches a height of 18.0 m above the pitcher's hand? (As you derive your answer, please show your work by typing equations, explanations, etc. in the box below.)

Please enter the answer to the following question in the space below. (Value: 6/100)

13. An object of mass M is dropped from 10.0 m above the water surface. Use conservation of mechanical energy to find the magnitude of its velocity 5.00 m above the water surface. Neglect air resistance and assume it starts from rest. (As you derive your answer, please show your work by typing equations, explanations, etc. in the box below.)

Please select one of the answers to the following question. (Value: 2/100)

14. A baseball with mass 0.145 kg is thrown straight up with an initial velocity of magnitude 25.0 m/s. If we ignore the air friction, when the baseball is 20.0 m above the ground, the magnitude of its velocity

1. does not depend on whether the baseball is moving upward or downward; correct; j
2. depend on whether the baseball is moving upward or downward; correct; j

Please enter the answer to the following question in the space below. (Value: 8/100)

15. A person pulls an 16.5 kg box across a floor by pulling on a rope with a constant force of 47.5 N. The rope is horizontal. The frictional force on the box is 42.0 N. Determine the magnitude of the velocity of the box after it has been pulled 6.50 m starting from an initial

velocity of 1.18 m/s. Use work energy concepts. Solutions using other methods will not be recognized. You may take the floor to be the zero level of gravitational potential energy. (As you derive your answer, please show your work by typing equations, explanations, etc. in the box below.)

Please select one of the answers to the following question.(Value: 2/100)

16. A block is slowly pushed against a spring, compressing it at a constant speed. During this time interval, the spring is exerting a force on the block. The work done by the spring on the block is:

1. positive
2. zero
3. negative

Please select one of the answers to the following question.(Value: 2/100)

17. Two men, Tom and Jerry, push against a wall with the same force. Jerry stops after 10 min, while Tom is able to push for 5 min longer. Compare the work against the wall they each do.

1. Tom does 50% more work than Jerry.
2. Jerry does 50% more work than Tom
3. Tom does 75% more work than Jerry.
4. Neither of them does any work.

Please select one of the answers to the following question.(Value: 2/100)

18. A person pulls a box along the rough ground with a constant magnitude of velocity. If we consider Earth and the box as our system, the work done by the person on the system is:

1. zero
2. nonzero

Please select one of the answers to the following question.(Value: 2/100)

19. When a student lifts a heavy box, the work done on the box by the Earth:

1. positive
2. negative
3. zero
4. need more information in order to answer.

Please select one of the answers to the following question.(Value: 2/100)

Q20*. An object can never have a negative kinetic energy. **Used in study 1**

1. True
2. False

Q20. For isolated systems, an increase in potential energy is always equal to: **Used in study 2 and 3**

1. an increase in kinetic energy
2. a decrease in kinetic energy
3. an increase in mechanical energy

4. a decrease in mechanical energy
5. cannot tell without more information.

Please select one of the answers to the following question.(Value: 6/100)

21. Suppose a heavy box is suspended by a rope. A man holding the rope slowly lowers the box onto a spring. The box is moving at constant velocity. Consider the system to consist of the box alone. The system's total mechanical energy:

1. increases.
2. decreases.
3. stays the same.
4. we cannot tell; more information is needed

Please select one of the answers to the following question.(Value: 6/100)

22. Suppose a heavy box is suspended by a rope. A man holding the rope slowly lowers the box onto a spring. The box is moving at constant velocity. Define the system to be the box and the spring, but no other objects. The system's total mechanical energy:

1. increases.
2. decreases.
3. stays the same.
4. we cannot tell; more information is needed.

Please select one of the answers to the following question.(Value: 6/100)

23. Suppose a heavy box is suspended by a rope. A man holding the rope slowly lowers the box onto a spring. The box is moving at constant velocity. Define the system to be the box and the earth but no other objects. The system's total mechanical energy:

1. increases.
2. decreases.
3. stays the same.
4. we cannot tell; more information is needed.

Please select one of the answers to the following question.(Value: 2/100)

24. You lift a ball at a constant velocity from a height h_i to a greater height h_f . Considering the ball ALONE to be the system, which of the following statements is true?

1. The potential energy of the system increases.
2. The kinetic energy of the system decreases.
3. The earth does negative work on the system.
4. You do negative work on the system.
5. The source energy of the ball increases.
6. Two of the above.
7. None of the above.

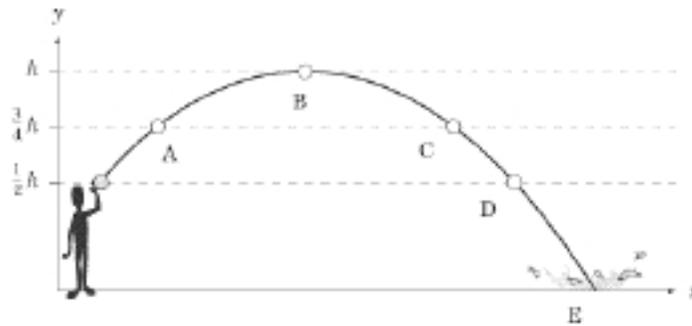
Please select one of the answers to the following question.(Value: 2/100)

25. You lift a ball at constant velocity from a height h_i to a greater height h_f . Considering the ball and the earth TOGETHER as the system, which of the following statements is true?

1. The potential energy of the system increases.

2. The kinetic energy of the system decreases.
3. The earth does negative work on the system.
4. You do negative work on the system.
5. The source energy of the ball increases.
6. Two of the above.
7. None of the above.

26. Consider the diagram of the trajectory of a thrown tomato: J



1. At what point is the potential energy greatest?
2. At what point is the kinetic energy the least?
3. At what point is the kinetic energy greatest?
4. At what point is the kinetic energy decreasing and the potential energy increasing?
5. At what two points are the kinetic energies equal and the potential energies equal?

Please select one of the answers to the following question.(Value: 2/100)

27. Two marbles, one twice as massive as the other, are dropped to the ground from the roof of a building. Ignore air resistance. Just before hitting the ground, the heavier marble has:

1. as much kinetic energy as the lighter one
2. twice as much kinetic energy as the lighter

3. one half as much kinetic energy as the lighter
4. one four times as much kinetic energy as the lighter one
5. square root of 2 (i.e., 1.414) times as much kinetic energy as the lighter one.
6. impossible to determine

Please select one of the answers to the following question.(Value: 2/100)

28. While flying along, a jet releases an empty propellant tank which eventually crashes to the ground. Consider three cases: a)the jet is flying horizontally, b) the jet is climbing upwards at a 45 degree angle, or, c) the jet is diving downwards at a 45 degree angle. In all 3 cases, it is flying at the same speed and it releases the propellant tank from the same height. Ignoring air friction, in which case is the tank's speed greatest as it hits the ground?

1. The horizontal jet
2. The climbing jet
3. The diving jet
4. It doesn't matter. The speed is the same in all three cases.
5. More information is needed in order to answer.

Please select one of the answers to the following question.(Value: 2/100)

29. Consider a system that consists of only the earth and box. A battery-powered motor not considered part of the system lifts the box from the ground to a certain height above the earth's surface. Which statement below is most accurate after the box has been lifted and come to a stop?

1. The kinetic energy of the box has increased.
2. The kinetic energy of the box has decreased.
3. The total mechanical energy of the system remains the same, because the box is not moving.

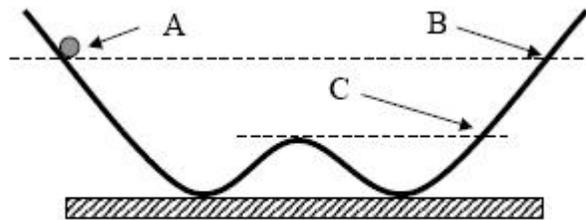
- The total mechanical energy of the system has decreased.
- The total mechanical energy of the system has increased.

Please select one of the answers to the following question.(Value: 2/100)

30. You support an object and move it to the right with a constant velocity. You exert a force F on it to oppose the gravitational attraction of the earth for the object. If you do not raise the object or increase its velocity, and air friction is negligible, do you do work on the object?

- Yes
- No

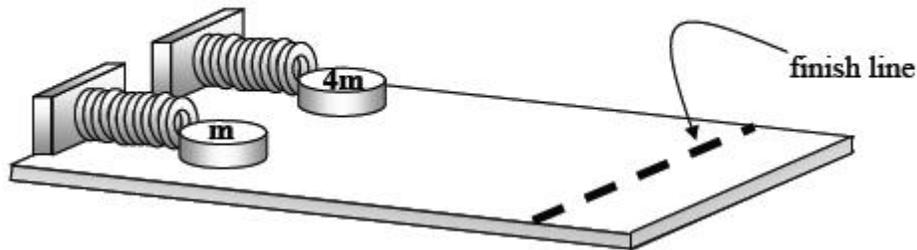
31. A steel ball is placed at position A on the curved, hard surface shown [insert diagram]. The surface is fixed to the table on which it sits so that it does not move. The ball is held at rest at position A and then is released. It rolls smoothly along the surface. Which choice below most nearly describes the greatest height attained by the ball on the other side of the curve?



- It will not get over the hill in the middle.
- Significantly below position C
- Almost to position C
- Almost to position B
- Slightly higher than position B

Please select one of the answers to the following question.(Value: 4/100)

32. Two pucks on level, frictionless ice are shown above, pressed back by equal amounts against identical springs. The pucks are the same size and shape, but one has four times the mass of the other. The pucks are released and the springs propel them to the finish line. At the finish line the kinetic energy of the less massive puck compare to the kinetic energy of the more massive puck?



1. the same as the kinetic energy of the more massive puck.
2. four times the kinetic energy of the more massive puck.
3. twice the kinetic energy of the more massive puck.
4. half the the kinetic energy of the more massive puck.
5. one-fourth the kinetic energy of the more massive puck.

33. A puck sitting on level ice is pushed back against a spring that is attached to a wall. This partially compresses the spring. The puck is released, and the spring propels it. If you have only this spring, but a variety of different pucks, how could another puck be given more energy? (i) Use a puck with less mass. (ii) Use a puck with more mass. (iii) Compress the spring more.

1. i only
2. ii only
3. iii only
4. i and iii
5. ii and iii

APPENDIX F

PRE- AND POST-TEST KCS

Table F1: KC Occurrences by Pre- & Post-test problem.

KC:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	Tot	
Problem 1										1																								1
2																					1													1
3																						1												1
4																																		1
5																																		1
6																																		1
7																																		1
8																																		1
9																																		3
10																																		2
11																																		2
12																																		6
13																																		10
14																																		11
15																																		15
16																																		1
17																																		1
18																																		1
19																																		1
20																																		2
21																																		2
22																																		3
23																																		3
24																																		5
25																																		5
26																																		8
27																																		7
28																																		7
29																																		11
30																																		11
31																																		11
32																																		14
33																																		14
Tot:	14	1	2	4	2	5	0	4	8	7	1	6	8	8	2	8	1	5	4	16	10	5	2	16	5	6	6	6	2	1	2	3	168	

APPENDIX G

TRAINING PROBLEMS

P1. At T_0 , a 2000 kg car is moving with a velocity of magnitude 22.35 m/s. Find the kinetic energy of the car.

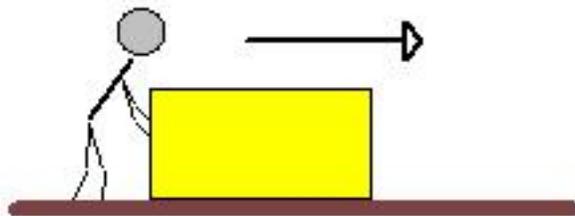
P2. A man pushes a 20.0 kg crate across a frictionless floor with a horizontal force of 24.0 N. What work is done by the man on the crate in displacing it by 5.0 m?

We define:

T_0 : the time point when the man starts pushing the crate.

T_1 : the time point when the man has finished pushing the crate 5.0 m

Please see diagram below.



P3. Jan's mountain bike has a spring with a constant of 64 N/m in the front wheel suspension. At T_0 the front wheel suspension gets compressed by 0.17m when she hits a bump. How much energy does the wheel-spring pair store at that point? Please see diagram

below.



P4. A 0.6 kg rock in space has a velocity of magnitude 2.0 m/s at point A and kinetic energy of 7.50 J at point B. What is the net work done on the rock as it moves from A to B?

We define

T0: the time point when the rock is at point A.

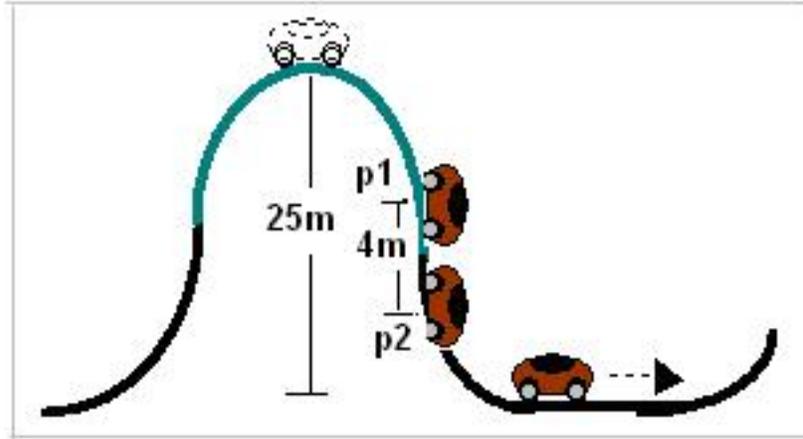
T1: the time point when the rock is at point B.

P5. Suppose a 2000kg truck is dropped from a certain height above the ground, starting from rest. Find the height from which it is dropped if it strikes the ground at 22.35 m/s. Neglect air resistance. Assume that the origin is at the ground. We define

T0: the time point when the truck is dropped, from rest.

T1: the time point when it strikes the ground.

P6. At time T0, a 225 kg frictionless roller-coaster car reaches the top of the first hill with a velocity of magnitude 3.0 m/s. At time T1, the car is at the point P1 and it reaches a velocity of magnitude 7.8 m/s. At time T2, the car is at the point P2, 4 meters lower in height than P1, and it reaches a velocity of magnitude 11.8 m/s. What is the magnitude of the car's instantaneous velocity at time T3, at the point when it reaches the base? Assume that the origin is at the base of the roller-coaster and the height of the first hill is 25 m above the base. Please see diagram below.



P7. A vertical spring is compressed by 0.15 m. A 3.00 kg block is then placed on the spring and released. At time T_1 , the block passes the spring's equilibrium position. Then the block flies into the air. At 0.2 s later, 0.404 m above the uncompressed spring, the block has a velocity of 1.04 m/s. How high will the block go above the uncompressed spring? What is the constant of the spring? Choose the origin to be at the top of the uncompressed spring. We define

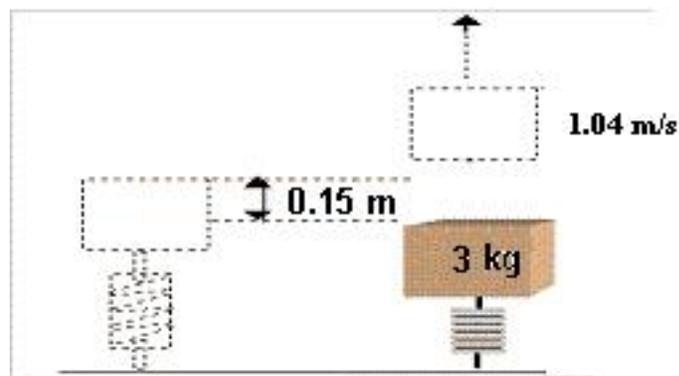
T_0 : the time point when the spring is compressed by 0.15 m with the block resting on it.

T_1 : the time point when the block passes the spring's equilibrium position.

T_2 : the time point when the block is 0.404 m above the spring's equilibrium position.

T_3 : the time point when the block reaches its maximum height.

Please see diagram below.



APPENDIX H

AN EXAMPLE OF STUDENT-CORDILLERA LOG FILE

Here is a student who is under the Exploratory Group solved P4 on Random-Cordillera.

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
1	14:51:41	Tutor:	[continue]This problem is a little more complex than the problems that you have solved so far, but I think you can handle it.	
2	14:51:41	Tutor:	[ent_txt]What does the problem statement ask you to find?	ELICIT
3	14:52:09	Student:	Work done on the rock	
4	14:52:26	Tutor:	[continue]It asks you to find net work, more specifically, net work done on the rock as it moves from A to B.	
5	14:52:36	Tutor:	[ent_txt]What are the units for net work?	ELICIT KC15
6	14:52:46	Student:	J	" +KC15"
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
7	14:52:47	Tutor:	Very good!	
8	14:52:47	Tutor:	[continue]Net work always refers to two time points.	
9	14:52:54	Tutor:	[continue]In this problem, they are T0 and T1. T0 is defined as the time when the rock is at point A and T1 is defined as the time when it is at point B.	
10	14:53:04	Tutor:	[continue]As always, let's first define variables and give values to some of them.	
11	14:53:09	Tutor:	[continue]I will define m to be the mass of the rock.	
12	14:53:14	Tutor:	[continue]There are two variables that we can define at both T0 and T1, the rock's magnitude of velocity and the rock's kinetic energy. I will define them for you. Look at the variable window after the next step.	
13	14:53:24	Tutor:	[ent_txt]Which variable has a value 0.60 kg?	ELICIT
14	14:53:46	Student:	M	
15	14:53:47	Tutor:	Yes.	
16	14:53:53	Tutor:	[ent_txt]Which variable has a value of 2.0 m/s?	ELICIT
17	14:54:16	Student:	V0	
18	14:54:17	Tutor:	Okay.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
19	14:54:18	Tutor:	[main_menu 'num_units'] What is the value of the kinetic energy of the rock at T1?	ELICIT KC23
20	14:54:46	Student:	1.2m/s	”-KC23”
21	14:54:47	Tutor:	I disagree with you.	
22	14:54:47	Tutor:	[continue]At T1,the rock is at point B, and the problem says the rock’s kinetic energy at point B is 7.50 J.	”*KC23”
23	14:54:52	Tutor:	[continue] Next let’s identify the sought variable.	
24	14:54:55	Tutor:	[continue]The problem asks for the net work done on the rock from A to B. I will define it as Wnet01 in the variable window.	TELL
25	14:55:00	Tutor:	[continue] Our main goal is to calculate Wnet01. There are two principles that are involved in net work. One is the definition of net work and the other is conservation of mechanical energy for non-isolated systems. Let’s examine whether one or both of these principles are applicable to this problem.	”*KC18 *KC28”
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
26	14:55:08	Tutor:	[continue]According to the definition of net work, the net work done on the rock is the sum of the work done by the external forces on the rock. From the problem statement, there is no information given about the forces acting on the rock(e.g. gravitational force). So we cannot apply the first principle directly to this problem.	”*KC18”
27	14:55:18	Tutor:	[continue]Now let’s examine the remaining principle involving net work, namely conservation of mechanical energy for non-isolated systems. It should be applicable since we don’t have any other principles to try!	
28	14:55:28	Tutor:	[continue]Conservation of mechanical energy for non-isolated systems in its generic form says that when there are external forces acting on a system, the net work done by these forces on the system equals the change in the total mechanical energy of the system.	”*KC28”
29	14:55:37	Tutor:	[continue]The first thing we need to do is to define a system because conservation of mechanical energy for non-isolated systems is applied to a system, not an object.	”*KC28”
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
30	14:55:45	Tutor:	[continue]In this problem, let's simply select the rock as the system.	TELL
31	14:55:48	Tutor:	justify_choose-system	SKIP- JUSTIFY
32	14:55:48	Tutor:	[continue]Next, we need to define variables for the total mechanical energy for the rock system at T0 and at T1. I will define them as TME0 and TME1 respectively. Please look at the variable window on the next step.	
33	14:55:56	Tutor:	[continue]There is more than one way to solve this problem. In the following discussion, I may lead you on a bit of a detour at times simply to help you practice applying some of the principles.	
34	14:56:03	Tutor:	[continue]We will solve this problem in three steps. First, we will calculate KE0, then we will calculate v1, and finally we will calculate Wnet01.	
35	14:56:08	Tutor:	[continue]So let's start with determining the value of KE0.	
36	14:56:11	Tutor:	[continue]To calculate the rock's kinetic energy at T0, let's apply the definition of kinetic energy.	TELL *KC20
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
37	14:56:27	Tutor:	[main_menu 'eqn']Please write the equation for applying the definition of kinetic energy to the rock at T0.	ELICIT KC20
38	14:56:31	Student:	$ke0 = 1/2 * M * V0^2$	" +KC20"
39	14:56:32	Tutor:	[main_menu 'num_units']Since all the variables in the equation are known except for KE0, we can calculate KE0 easily. What is the value of KE0?	ELICIT KC23
40	14:56:47	Student:	1.2J	" +KC23"
41	14:56:48	Tutor:	Okay.	
42	14:56:49	Tutor:	[continue]Before we go any further, let's discuss this principle in a little more detail.	
43	14:56:56	Tutor:	[ent_txt]Why does the rock have a non-zero kinetic energy at T0?	Justify ELICIT KC20
44	14:57:17	Student:	because it has a nonzero mass and a nonzero magnitude of velocity	" +KC20"
45	14:57:18	Tutor:	Yes.	
46	14:57:19	Tutor:	[continue]We do not know the direction of the rock's velocity at T0 from the problem statement. However, the direction of v0 does not affect the rock's kinetic energy at T0.	TELL *KC20
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
47	14:57:23	Tutor:	58-ke0_discuss_qualitative_2	Skip- JUSTIFY KC20
48	14:57:23	Tutor:	[continue]We've calculated KE0, recall that our plan is to calculate v1 next.	
49	14:57:26	Tutor:	[continue]To calculate the rock's instantaneous magnitude of velocity at T1, we will apply the definition of kinetic energy again.	TELL *KC20
50	14:57:31	Tutor:	[continue]Let me just write the equation for you: $KE1 = 1/2 * m * v1^2$.	TELL *KC20
51	14:57:34	Tutor:	[continue]From $KE1 = 1/2 * m * v1^2$, we get $v1^2 = KE1 / (0.5 * m)$. We substitute KE1 with 7.50 J and m with 0.60 kg. The equation now becomes: $v1^2 = 7.50 \text{ J} / (0.5 * 0.60 \text{ kg})$. The magnitude of velocity is always positive, therefore, $v1 = 5.0 \text{ m/s}$.	TELL *KC31
52	14:57:41	Tutor:	58-v1_discuss_qualitative	SKIP- Justify KC20
53	14:57:42	Tutor:	[continue]Now that we know v1, we will concentrate on the sought quantity for this problem, that is, the net work done on the rock system from A to B.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
54	14:57:48	Tutor:	[continue]Our system consists of the rock and we have planned to apply the conservation of mechanical energy for non-isolated systems from T0 to T1 to find the net work done on the rock from point A to point B.	”*KC28”
55	14:57:54	Tutor:	[continue]More specifically, we will first find the system’s total mechanical energy at T0, then we will find it at T1, and finally we will find the net work done on the rock from A to B by examining the change in total mechanical energy from T0 to T1.	
56	14:58:00	Tutor:	[continue]Let’s use the definition of total mechanical energy to find TME0.	TELL *KC24
57	14:58:04	Tutor:	[continue]The generic equation for the definition of total mechanical energy is $TME = KE + GPE + SPE$.	TELL *KC24
58	14:58:08	Tutor:	[continue]In this problem, the system only has one object, the rock. Therefore, there are no potential energies involved, neither gravitational potential energy nor spring potential energy.	”*KC21 *KC22”
59	14:58:09	Tutor:	58-tme0_discuss_truth	SKIP- JUSTIFY KC21
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
60	14:58:10	Tutor:	[main_menu 'eqn'] Please write the equation for the definition of total mechanical energy for the rock system at T0.	ELICIT KC24
61	14:58:33	Student:	$TME0 = KE0$	" +KC24"
62		Tutor:	[main_menu 'num_units']It is easy to infer the value of TME0. Please calculate the value of TME0.	ELICIT KC23
63	14:58:52	Student:	1.2J	" +KC23"
64	14:58:53	Tutor:	[cust_menu 'Yes, it is true—No, it is not true—I do not know—Not enough information']Equation $TME0 = KE0$ tells us that the rock system's total mechanical energy is equal to the rock's kinetic energy at T0. Is it true *for any given time*?	ELICIT KC24
65	14:59:29	Student:	No, it is not true	" -KC24"
66	14:59:30	Tutor:	I disagree with you.	
67	14:59:31	Tutor:	[continue]In factit is true for any time point in this problem.	"*KC24"
68	14:59:32	Tutor:	[continue]This is because as long as the rock is the only object in the system, there are no potential energy involved, neither gravitational potential energy nor spring potential energy. Thus $TME = KE$ at any given time point.	TELL *KC24 *KC21 *KC22
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
69	14:59:40	Tutor:	58-tme0_discuss_relevance	SKIP- JUSTIFY KC28
70	14:59:41	Tutor:	[continue]We have computed TME0, let's calculate TME1 next.	
71	14:59:43	Tutor:	[continue]Since the total mechanical energy is equal to kinetic energy at any given time as long as the rock is the only object in a system, let's use the definition of total mechanical energy again to get the value for TME1.	TELL *KC24
72	14:59:50	Tutor:	[main_menu 'eqn'] Please write the appropriate equation for the definition of total mechanical energy of the system at T1.	ELICIT KC24
73	14:59:58	Student:	$TME1 = KE1$	" +KC24"
74	14:59:59	Tutor:	[main_menu 'num_units']What is the value of TME1?	ELICIT *KC23
75	15:00:16	Student:	$7.5J$	
76	15:00:17	Tutor:	58-tme1_discuss_qualitative	SKIP- JUSTIFY KC24
77	15:00:17	Tutor:	58-tme1_discuss_truth	SKIP- JUSTIFY KC24
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
78	15:00:17	Tutor:	[continue]We just calculated TME1 and TME0, let's calculate the net work done on the rock from A to B.	
79	15:00:22	Tutor:	[continue]As we have mentioned before, there were two principles from the textbook which apply to net work: one is the definition of net work and the other is conservation of mechanical energy for non-isolated systems.	TELL *KC28 *KC18
80	15:00:30	Tutor:	[continue]The definition of the net work cannot be applied to this problem because there is no information given in the problem statement about the forces on the rock. Therefore, we will apply the remaining principle, namely, conservation of mechanical energy for non-isolated systems. We will start with the generic equation.	
81	15:00:38	Tutor:	[main_menu 'form'] What is the generic equation for conservation of mechanical energy for non-isolated systems?	ELICIT KC28q
82	15:01:06	Student:	$TME0 = TME1$	"-KC28q"
83	15:01:07	Tutor:	That doesn't sound right to me.	
84	15:01:08	Tutor:	[continue]The generic equation for conservation of mechanical energy for non-isolated systems: $W_{net01} = TME1 - TME0$	"*KC28q"
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
85	15:01:15	Tutor:	[ent_txt]What time points and bodies should we use when applying the principle of conservation of mechanical energy for non-isolated systems?	ELICIT KC28
86	15:01:57	Student:	We should use the time points where the total mechanical energy changes.	”-KC28”
87	15:01:59	Tutor:	[ent_txt]What objects should you apply the conservation of mechanical energy for non-isolated systems to?	
88	15:02:29	Student:	the objects that are in motion	”-KC28”
89	15:02:31	Tutor:	[continue]We have already defined the system to be the rock. Therefore, we should apply conservation of mechanical energy to the rock system.	”*KC28”
90	15:02:37	Tutor:	[continue]To apply conservation of mechanical energy for non-isolated systems, we need two time points.	
91	15:02:47	Tutor:	[ent_txt]What time points should we use when applying conservation of mechanical energy for non-isolated systems?	
92	15:03:30	Student:	T0 and T1	” +KC28”
93	15:03:32	Tutor:	Okay.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
94	15:03:32	Tutor:	[continue]Let me just write the equation for conservation of mechanical energy for the rock system between T0 and T1 for you: $TME0 + W_{net01} = TME1$	TELL *KC28
95	15:03:39	Tutor:	[continue]Since we already know the values for TME0 and TME1, we can calculate the value of Wnet01 easily.	ELICIT KC15
96	15:03:46	Tutor:	[main_menu 'num_units']Please calculate the value of Wnet01.	
97	15:03:56	Student:	6.3J	” +KC15”
98	15:03:58	Tutor:	[continue]We have finished solving the whole problem. Next let’s do a little discussion.	
99	15:04:00	Tutor:	[cust_menu ‘KE0—v1—TME1—TME0—Wnet01’]Are any of these variables necessary—I do not know. [While solving the problem, we calculated the variables in the order of KE0, v1, TME0, TME1, and Wnet01. Are any of these variable(s) not necessary to solve the problem?	ELICIT
100	15:04:54	Student:	v1	
101	15:04:55	Tutor:	Yes.	
102	15:04:57	Tutor:	[continue]The reason that we calculate v1 in this problem is to practice the principle of the definition of kinetic energy.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
103	15:05:02	Tutor:	[cust_menu ‘Yes, it is possible—No, it is not possible—I do not know—Not enough information’]In this problem, we have selected the rock as the system. Is it possible to define potential energy for the rock system?	ELICIT KC21
104	15:05:22	Student:	No, it is not possible	”+KC21”
105	15:05:23	Tutor:	Yes, I agree.	
106	15:05:23	Tutor:	58-pe-concept_justification_2	SKIP- Justify KC21
107	15:05:23	Tutor:	[continue]Next, we will change the place where the rock is and discuss whether the values of some variables would also change.	
108	15:05:31	Tutor:	[cust_menu ‘increase—decrease—same—not enough information—I do not know’]Suppose the 0.6 kg rock is put on the Earth and the rock’s kinetic energy at point B is still 7.5 J, what will happen to the rock’s magnitude of velocity at the point B?	ELICIT KC20
109	15:06:09	Student:	same	”+KC20”
110	15:06:10	Tutor:	Very good!	
111	15:06:10	Tutor:	58-v1-ppd-discussion-2	SKIP- JUSTIFY KC20
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
112	15:06:10	Tutor:	[cust_menu ‘increase—decrease—same—not enough information—I do not know.’]Suppose the rock is on the Earth and everything in the problem statement is kept the same. What will happen to net work done on the rock if we still choose the rock alone as the system?	ELICIT KC28 KC24
113	15:06:36	Student:	same	”+KC28 +KC24”
114	15:06:37	Tutor:	Sounds good.	
115	15:06:37	Tutor:	58-wnet01-ppd-discussion-2	SKIP- JUSTIFY KC28 KC24
116	15:06:37	Tutor:	[main_menu ‘num_units’]Let’s go back to the situation where rock is in space. Look at the variable window, tell me what is the magnitude of v_0 ?	
117	15:06:55	Student:	$2m/s$	”+KC31”
118	15:06:56	Tutor:	Yes, I agree.	
119	15:06:57	Tutor:	[main_menu ‘num_units’]What is the magnitude of v_1 ?	
120	15:07:08	Student:	$5m/s$	”+KC31”
121	15:07:09	Tutor:	Excellent!	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
122	15:07:11	Tutor:	[continue]Although the magnitude of v_1 (5.0 m/s) is bigger than the magnitude of v_0 (2.0 m/s), we *cannot* infer that the rock's magnitude of velocity is increasing during the ENTIRE period from T_0 to T_1 .	TELL *KC31 *KC20
123	15:07:17	Tutor:	[main_menu 'num_units']Look at the variable window, tell me what is the value of the rock's kinetic energy at T_0 ?	
124	15:07:31	Student:	1.2J	" +KC23"
125	15:07:32	Tutor:	Yes, I agree.	
126	15:07:34	Tutor:	[main_menu 'num_units']What is the value of KE1?	
127	15:07:47	Student:	7.5J	" +KC23"
128	15:07:48	Tutor:	Excellent!	
129	15:07:50	Tutor:	[continue]Compared with point A, the kinetic energy of the rock at point B is increased.	TELL *KC20
130	15:07:52	Tutor:	[ent_txt]Please try to explain why.	Justify ELICIT KC20
131	15:08:17	Student:	it is so because the velocity of the rock has increased from A to B	" +KC20"
132	15:08:18	Tutor:	Very good!	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
133	15:08:19	Tutor:	[continue]Suppose there is point C on the path of the rock between A to B. We do not have enough information to say anything about the kinetic energy of the rock at point C compared with point A.	TELL *KC20
134	15:08:27	Tutor:	[continue]This is because we do not know whether the rock's magnitude of velocity at C is bigger than its magnitude of velocity at A or not. It is possible that the rock continuously increases its magnitude of velocity from A to C then to B. It is also possible its magnitude of velocity decreases from A to C and then increases from C to B. It is even possible that it does not change its magnitude of velocity from A to C and then increases its magnitude of velocity from C to B.	Justify TELL *KC20
135	15:08:39	Tutor:	[menu 'increase—decrease—same—not enough information—I do not know.']Compared with point B, What can you say about the kinetic energy of the rock at point C?	ELICIT KC20
136	15:09:12	Student:	not enough information	" +KC20"
137	15:09:13	Tutor:	Sounds good.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
138	15:09:13	Tutor:	[ent_txt]Please explain why.	JUSTIFY ELICIT KC20
139	15:10:07	Student:	It is so because the KE depends on the velocity and we do not know the velocity of the rock at point C.	”+KC20”
140	15:10:08	Tutor:	[continue]Next, we will change the mass of the rock and discuss whether the values of some variables would also change.	
141	15:10:10	Tutor:	[continue]Suppose that the rock’s mass increased from 0.6 kg to 1 kg and the rock’s kinetic energy at point B was still 7.5 J, the rock’s magnitude of velocity at point B would decrease.	TELL *KC20
142	15:10:16	Tutor:	[continue]It is because we are assuming that the rock’s kinetic energy at point B stays same, 7.50 J, and if the rock’s mass increases, according to $KE = \frac{1}{2} * m * v^2$, then the magnitude of rock’s magnitude of velocity has to decrease.	JUSTIFY TELL *KC20
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
143	15:10:22	Tutor:	[cust_menu ‘increase—decrease—same—not enough information—I do not know.’]Compared with point B, What can you say about the kinetic energy of the rock at point C?	ELICIT KC20
144	15:10:46	Student:	increase	”+KC20”
145	15:10:47	Tutor:	I agree.	
146	15:10:49	Tutor:	[ent_txt]Why?	JUSTIFY ELICIT KC20 KC23
147	15:11:09	Student:	It is so because the KE of a body depends on it’s mass	”-KC20”
148	15:11:10	Tutor:	I disagree with you.	
149	15:11:12	Tutor:	[continue]Because the magnitude of rock magnitude of velocity is still 2.0 m/s at point A and the mass of the rock increases, according to $KE = \frac{1}{2} * m * v^2$ the kinetic energy of the rock at point A increases.	*KC20 *KC20q *KC31
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
150	15:11:18	Tutor:	[cust_menu ‘increase—decrease—same—not enough information—I do not know.’]Finally, if the rock’s mass were 1 kg instead of 0.6 kg and the rock’s magnitude of velocity at point A was still 2.0 m/s and the rock’s kinetic energy at point B was still 7.5 J, what will happen to the net work done on the rock as it moves from A to B?	ELICIT KC28 KC24 KC20
151	15:11:59	Student:	decrease	”+KC28, +KC24, +KC20”
152	15:12:00	Tutor:	58-inc_mass3.2	SKIP- JUSTIFY KC28 KC24 KC20
153	15:12:02	Tutor:	[cust_menu ‘Yes, we can—No, we cannot—not enough information—I do not know.’]If we already know that the work done on the rock by force F is 9 J and we also know that the rock is moving in the same direction as F for 3 m in a straight line, can you infer the magnitude of the force F?	ELICIT KC14
154	15:12:28	Student:	Yes, we can	”-KC14”
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
155	15:12:29	Tutor:	I disagree with you.	
156	15:12:31	Tutor:	[continue]No, we cannot infer anything about the magnitude of the force F.	
157	15:12:39	Tutor:	58-work_concept_justification	SKIP- JUSTIFY KC14
158	15:12:39	Tutor:	[continue]Let's go back to the original situation. The problem does not give any hint about the path taken by the rock. However, taking different paths would not affect our final answer of the net work done on the rock from A to B.	TELL *KC28 *KC24 *KC20
159	15:12:49	Tutor:	58-compare_path_2	SKIP- JUSTIFY KC28 KC24 KC20 KC15
160	15:12:49	Tutor:	[continue]We're finished with this problem. I think I should give you a more complex problem next time.	
Continued on Next Page...				

Table H1: An Example of Student Log on Training Problem P4

	Time	speaker	string	Label
161	15:12:59	Tutor:	[ent_txt]The next problem will take about 40 minutes to complete. If you start this problem and do not have enough time to finish, then you will have to start it over from the beginning when you return for your next session. Are you ready for the next problem?	
162	15:13:30	Student:	Yes	

APPENDIX I

AN EXAMPLE TUTORIAL SCRIPT

Example tutorial script for problem 58, listed as problem 4 in Appendix [G](#).

```
g start
  say ‘‘[start_problem 58]’’
  do solving_problem
  do discuss_problem
  say ‘‘[cont]We’re finished with this problem. I think I should give
      you a more complex problem next time.’’
  do logout

g solving_problem
  say ‘‘[cont]This problem is a little more complex than the problems
      that you have solved so far, but I think you can handle
      it.’’
  do pre-discuss-problem
  do time_points
  say ‘‘[cont]As always, let’s first define variables and give values
      to some of them.’’
```

```

do given_soughts
do application

g pre-discuss-problem sem elicit
say ‘‘What does the problem statement ask you to find?’’
if ‘‘net work.’’ true
if ‘‘net work done on the rock as it moves from A to B.’’ true
otherwise say ‘‘[cont]It asks you to find net work, more specifically,
        net work done on the rock as it moves from A to B.’’
say ‘‘What are the units for net work?’’
if ‘‘Joules.’’ true kc ‘‘+KC15’’
otherwise say ‘‘[cont]The units of the net work are Joules.’’
        kc ‘‘-KC15 *KC15’’

g pre-discuss-problem sem tell
say ‘‘[cont]The problem statement asks us to find the net work done
on the rock when it moves from point A to point B.’’
say ‘‘[cont]The units for net work are Joules. ’’ kc ‘‘*KC15’’

g time_points
say ‘‘[cont]Net work always refers to two time points.’’
say ‘‘[cont]In this problem, they are T0 and T1. T0 is defined as the
time when the rock is at point A and T1 is defined as
        the time when it is at point B.’’
say ‘‘[add 'T0']’’
say ‘‘[add 'T1']’’
say ‘‘[add 'rock']’’

g given_soughts

```

```

do list_givens
do list_soughts

g list_givens
  say ''[cont]I will define m to be the mass of the rock.''
  say ''[add 'm' omit='value']''
  say ''[cont]There are two variables that we can define at both
      T0 and T1, the rock's magnitude of velocity and the
      rock's kinetic energy. I will define them for you.
      Look at the variable window after the next step.''
  say ''[add 'v0' omit='mag']''
  say ''[add 'KE0' omit='value']''
  say ''[add 'v1' omit='mag']''
  say ''[add 'KE1' omit='value']''
  do list_mass
  do list_v0
  do list_ke1

g list_mass sem elicit
  say ''Which variable has a value {m.value}?''
  if ''m'' true
  otherwise say ''[cont]It is m.''
  say ''[add 'm']''

g list_mass sem tell
  say ''[cont]From the problem statement, we know that m
      has a value {m.value}.''
  say ''[add 'm']''

```

```

g list_v0 sem elicit
say ''Which variable has a value of 2.0 m/s?''
if ''v0'' true
otherwise say ''[cont]v0. ''
say ''[add 'v0']''

g list_v0 sem tell
say ''[cont]We can also see from the problem statement that
          v0 has a value of 2.0 m/s.'' kc ''*KC31''
say ''[add 'v0']''

g list_ke1 sem elicit
say ''[ask 'num_units'] What is the value of the kinetic
          energy of the rock at T1?''
if ''{KE1.value}'' true kc ''+KC23''
if ''correct value, wrong/no unit''
    say ''[cont]You entered
          the correct value, but not the correct units.
          The correct units for kinetic energy are J. ''
    kc ''-KC23 *KC23''
if ''{KE0.value}''
    say ''[cont] You seem to have calculated the value for
          the kinetic energy of the rock at T0. Now we
          are still listing the given values. At T1,the
          rock is at point B, and the problem says the
          rock's kinetic energy at point B is {KE1.value}.
          '' kc ''+KC23''
otherwise say ''[cont]At T1,the rock is at point B, and the
          problem says the rock's kinetic energy at
          point B is {KE1.value}.'' kc ''*KC23''

```

```

say ''[add 'KE1']''

g list_ke1 sem tell
say ''[cont]At T1,the rock is at point B, and the problem
           says the rock's kinetic energy at point B
           is {KE1.value}'' kc ''*KC23''
say ''[add 'KE1']''

g list_soughts
  say ''[cont] Next let's identify the sought variable.''
  do list_h0

g list_h0 sem elicit
say ''What is the sought quantity for this problem?''
if ''net work done on the rock from T0 to T1.'' true
  say ''[cont]I defined it as Wnet01 in the
           variable window.''
otherwise say ''[cont]The problem asks for the net work
           done on the rock from A to B. I defined
           it as Wnet01 in the variable window.''
say ''[add 'Wnet01' omit='value']''

g list_h0 sem tell
say ''[cont]The problem asks for the net work done on the rock
           from A to B. I will define it as Wnet01 in the
           variable window.''
say ''[add 'Wnet01' omit='value']''

g application

```

```

do global_discuss
say ''[cont]So let's start with determining the value of KE0.''
do KE0
say ''[cont]We've calculated KE0, recall that our plan is to
    calculate v1 next.''
do v1
say ''[cont]Now that we know v1, we will concentrate on the sought
    quantity for this problem, that is, the net work done on
    the rock system from A to B. ''
say ''[cont]Our system consists of the rock and we have planned to
    apply the conservation of mechanical energy for
    non-isolated systems from T0 to T1 to find the net work done
    on the rock from point A to point B.'' kc ''*KC28''
say ''[cont]More specifically, we will first find the system's total
    mechanical energy at T0, then we will find it at T1, and
    finally we will find the net work done on the rock
    from A to B by examining the change in total mechanical
energy from T0 to T1.''
do TME0
say ''[cont]We have computed TME0, let's calculate TME1 next.''
do TME1
say ''[cont]We just calculated TME1 and TME0, let's calculate the net
    work done on the rock from A to B.''
do CME01

g global_discuss
say ''[cont] Our main goal is to calculate Wnet01. There are two
    principles that are involved in net work. One is the definition of
    net work and the other is conservation of mechanical energy for non-
    isolated systems. Let's examine whether one or both of these prin

```

principles are applicable to this problem.’’ kc ’’*KC18 *KC28’’

say ’’[cont]According to the definition of net work, the net work done on the rock is the sum of the work done by the external forces on the rock. From the problem statement, there is no information given about the forces acting on the rock(e.g. gravitational force). So we cannot apply the first principle directly to this problem.’’ kc ’’*KC18’’

say ’’[cont]Now let’s examine the remaining principle involving net work, namely conservation of mechanical energy for non-isolated systems. It should be applicable since we don’t have any other principles to try!’’

say ’’[cont]Conservation of mechanical energy for non-isolated systems in its generic form says that when there are external forces acting on a system, the net work done by these forces on the system equals the change in the total mechanical energy of the system.’’ kc ’’*KC28’’

say ’’[cont]The first thing we need to do is to define a system because conservation of mechanical energy for non-isolated systems is applied to a system, not an object.’’ kc ’’*KC28’’

do choose-system_1

say ’’[cont]Next, we need to define variables for the total mechanical energy for the rock system at T0 and at T1. I will define them as TME0 and TME1 respectively. Please look at the variable window on the next step. ’’

say ’’[add ’TME0’ omit=’value’]’’

say ’’[add ’TME1’ omit=’value’]’’

say ’’[cont]There is more than one way to solve this problem. In the following discussion, I may lead you on a bit of a detour at times simply to help you practice applying some of the principles. ’’

say ’’[cont]We will solve this problem in three steps. First, we will

calculate KE0, then we will calculate v1, and finally we will
calculate Wnet01.''

```
g choose-system_1
  do choose-system
  do justify_choose-system sem justify
```

```
g choose-system sem elicit
  say ''[add 'rock']''
  say ''What would be your choice of the system for this problem?''
  if ''rock'' true
  otherwise do bottom_out_choose-system
```

```
g justify_choose-system sem elicit
  say ''Why?''
  if ''In this problem there is only one object: the rock. It is the
  target object.'' true
  otherwise say ''[cont]In this problem the target object is the rock
  . Since we have no information about the forces being applied to
  it, we do not have any clue whether there exist any other objects
  besides the rock. Therefore, we choose the rock alone to be the
  system.''
```

```
g justify_choose-system sem tell
  say ''[cont]In this problem the target object is the rock. Since we
  have no information about the forces being applied to it, we do
  not have any clue whether there exist any other objects besides
  the rock. Therefore, we choose the rock alone to be the system.''
```

```

    g bottom_out_choose-system
say ''[cont]Let's choose the rock to be the system.''

g choose-system sem tell
    say ''[add 'rock']''
    say ''[cont]In this problem, let's simply select the rock as the
    system. ''

g KEO
do KEO_principle_selection
do KEO_write_eqn
do KEO_solve_equation
do KEO_discuss_truth rand sem justify
do KEO_discuss_qualitative

g KEO_principle_selection sem elicit
    say ''Which principle will help you calculate the rock's kinetic
    energy at T0? Please provide the name of the principle, not
    an equation. ''
    if ''Definition of kinetic energy. '' true kc ''+KC20''
    otherwise say ''[cont]Let's apply the definition of kinetic
    energy.'' kc ''-KC20 *KC20''
    say ''[ask 'form'] What is the generic equation for the definition
    of kinetic energy? KE=''
    if ''{form_KE}'' true kc ''+KC20q''
    otherwise say ''[cont]The generic equation for the definition of

```

```

    kinetic energy is:  $KE = \frac{1}{2}mv^2$ '' kc ''-KC20q *KC20q''

g KEO_principle_selection sem tell
say ''[cont]To calculate the rock's kinetic energy at T0, let's
  apply the definition of kinetic energy. '' kc ''*KC20''
say ''[cont]The generic equation for the definition of kinetic
  energy is  $KE = \frac{1}{2}mv^2$ .'' kc ''*KC20q''

g KEO_write_eqn sem elicit
say ''[ask 'eqn']Please write the equation for applying the
  definition of kinetic energy to the rock at T0.''
if ''{eqn_KEO}'' true say ''[add 'eqn_KEO']'' kc ''+KC20q''
if ''equation with substituted values'' do bo_KEO_sub_write_eqn
otherwise do bo_KEO_write_eqn kc ''-KC20q''

g KEO_write_eqn sem tell
say ''[cont]Now I will write the equation for applying the
  definition of kinetic energy to the rock at T0:  $\frac{1}{2}mv^2$ ''
  kc ''*KC20q''
say ''[add 'eqn_KEO']''

g bo_KEO_sub_write_eqn
say ''[cont]You substituted problem values into the equation.
  Please use the variable labels instead.''
do KEO_write_eqn

g bo_KEO_write_eqn
say ''[add 'eqn_KEO']''
say ''[cont]Let me write the equation for the rock's kinetic

```

energy at T0 for you: {eqn_KE0}. Please look at the equation window.'' kc ''*KC20q''

g KE0_solve_equation sem elicit

say ''[ask 'num_units']Since all the variables in the equation are known except for KE0, we can calculate KE0 easily. What is the value of KE0?''

if ''{KE0.value}'' true kc ''+KC23''

if ''correct value, wrong/no unit'' say ''[cont]You entered the correct value, but not the correct units. The correct units for kinetic energy are J. '' kc ''-KC23 *KC23''

otherwise do bo_KE0_solve_equation

say ''[add 'KE0']''

g bo_KE0_solve_equation

say ''[cont]I calculate the value for $KE0 = 0.5 * \{m.value\} * \{v0.mag\}^2 = \{KE0.value\}$ '' kc ''*KC23''

g KE0_solve_equation sem tell

say ''[cont]Since all the variables in the equation are known except KE0, we can calculate KE0 easily. $KE0 = 0.5 * \{m.value\} * \{v0.mag\}^2 = \{KE0.value\}$ '' kc ''*KC23''

say ''[add 'KE0']''

g KE0_discuss_truth sem elicit

say ''[cont]Before we go any further, let's discuss this principle in a little more detail.''

say ''Why does the rock have a non-zero kinetic energy at T0?''

if ''Because the rock has a non-zero velocity at T0.''' true kc ''

```

+KC20''
if ''Because the rock is moving at T0'' true kc ''+KC20''
if ''Because the rock's magnitude of velocity at T0 is above 0''
true kc ''+KC20''
otherwise say ''[cont]Because the rock has a non-zero velocity
at T0 and the rock's mass is not zero.'' kc ''-KC20 *KC20''

g KEO_discuss_truth sem tell
say ''[cont]Because the rock has a non-zero velocity at T0 and
the rock's mass is not zero, the rock has a non-zero kinetic
energy at T0.'' kc ''*KC20''

g KEO_discuss_qualitative
do KEO_discuss_qualitative_1
do KEO_discuss_qualitative_2 rand sem justify

g KEO_discuss_qualitative_1 sem elicit
say ''[menu 'Yes, it would affect|No, it would not affect|I do
not know|Not enough information']We do not know the direction
of the rock's velocity at T0 from the problem statement. Would
the direction of v0 affect the rock's kinetic energy at T0?''
if ''No, it would not affect'' true kc ''+KC20''
otherwise say ''[cont]Actually ... the direction of the rock at
T0 would not affect the rock's kinetic energy at T0.''
kc ''-KC20 *KC20''

g KEO_discuss_qualitative_1 sem tell
say ''[cont]We do not know the direction of the rock's velocity
at T0 from the problem statement. However, the direction of v0

```

```

    does not affect the rock's kinetic energy at T0.'' kc ''*KC20''

g KE0_discuss_qualitative_2 sem elicit
  say ''Why?''
  if ''The rock's kinetic energy only depends on the rock's mass
    and its magnitude of velocity, not the direction of the rock.''
    true kc ''+KC20''
  if ''The direction of the rock at T0 does not matter. Only
    magnitude of velocity matters.'' true kc ''+KC20''
  otherwise say ''[cont]The direction of the rock at T0 would not
    affect the rock's kinetic energy at T0 because the kinetic
    energy depends on mass and the magnitude of velocity, not the
    direction of the rock's velocity.'' kc ''-KC20 *KC20''

g KE0_discuss_qualitative_2 sem tell
  say ''[cont]This is because the kinetic energy only depends on
    the rock's mass and magnitude of velocity, not the direction of
    the rock's velocity.'' kc ''*KC20''

g v1
  do v1_principle_selection
  do v1_write_eqn
  do v1_solve_equation
  do v1_discuss_qualitative rand
  .

g v1_principle_selection sem elicit

```

```

say ''Which principle will help you calculate the rock's
instantaneous magnitude of velocity at T1?''
if ''Definition of kinetic energy.'' true kc ''+KC20''
otherwise say ''[cont]We should apply the definition of kinetic
energy once again.'' kc ''-KC20 *KC20''

g v1_principle_selection sem tell
say ''[cont]To calculate the rock's instantaneous magnitude of
velocity at T1, we will apply the definition of kinetic energy
again. '' kc ''*KC20''

g v1_write_eqn sem elicit
say ''[ask 'eqn']Please write the equation for how the
definition of kinetic energy applies to this problem at T1.''
if ''{eqn_KE1}'' true say ''[add 'eqn_KE1']'' kc ''+KC20q''
if ''equation with substituted values'' do bo_v1_sub_write_eqn
otherwise do bo_v1_write_eqn kc ''-KC20q''

g bo_v1_sub_write_eqn
say ''[cont]You substituted problem values into the equation.
Please use the variable labels instead.''
do v1_write_eqn

g bo_v1_write_eqn
say ''[add 'eqn_KE1']''
say ''[cont]Let me just write the equation for you: {eqn_KE1.
value}.'' kc ''*KC20q''

g v1_write_eqn sem tell
say ''[add 'eqn_KE1']''

```

```
say ''[cont]Let me just write the equation for you: {eqn_KE1.value}.'' kc ''*KC20q''
```

```
g v1_solve_equation sem elicit
```

```
say ''[ask 'num_units']Now it is easy to calculate the magnitude of v1. What is the magnitude of v1?''
```

```
if ''{v1.mag}'' true say ''[add 'v1']'' kc ''+KC31''
```

```
if ''correct value, wrong/no unit'' say ''[cont]You entered the correct value, but not the correct units. The correct units for velocity are m/s. '' kc ''-KC31 *KC31''
```

```
otherwise do bo_v1_solve_equation
```

```
say ''[add 'v1']''
```

```
g bo_v1_solve_equation
```

```
say ''[cont]From {eqn_KE1.value}, we get  $v1^2=KE1/(0.5*m)$ . We substitute KE1 with {KE1.value} and m with {m.value}. The equation now becomes:  $v1^2={KE1.value}/(0.5*{m.value})$ . The magnitude of velocity is always positive, therefore,  $v1= {v1.mag}$ . '' kc ''*KC31''
```

```
g v1_solve_equation sem tell
```

```
say ''[add 'v1']''
```

```
say ''[cont]From {eqn_KE1.value}, we get  $v1^2=KE1/(0.5*m)$ . We substitute KE1 with {KE1.value} and m with {m.value}. The equation now becomes:  $v1^2={KE1.value}/(0.5*{m.value})$ . The magnitude of velocity is always positive, therefore,  $v1= {v1.mag}$ .'' kc ''*KC31''
```

```

g v1_discuss_qualitative
  say ''[cont]Before going on to the next step, let's think about
    the application of this equation.''
  do v1_discuss_qualitative1
  do v1_discuss_qualitative2 sem justify

g v1_discuss_qualitative1 sem elicit
  say ''[menu 'Yes, we can|No, we cannot|I do not know']Can we
    infer the direction of the velocity of the rock at T1 from the
    rock's kinetic energy at T1?''
  if ''No. we cannot.'' true kc ''+KC20''
  otherwise say ''[cont]No. We cannot get the direction of the
    rock's velocity.'' kc ''-KC20 *KC20''

g v1_discuss_qualitative1 sem tell
  say ''[cont]We cannot infer anything about the direction of the
    rock at T1 from the rock's kinetic energy at T1. '' kc ''*KC20''

g v1_discuss_qualitative2 sem elicit
  say ''Please explain why.''
  if ''Because the kinetic energy only depends on mass and the
    magnitude of velocity, not the direction of velocity.''
    true kc ''+KC20''
  otherwise say ''[cont]This is because the kinetic energy only
    depends on mass and the magnitude of velocity, not the
    direction of velocity. '' kc ''-KC20 *KC20''

g v1_discuss_qualitative2 sem tell
  say ''[cont]This is because the kinetic energy only depends

```

on mass and the magnitude of velocity, not the direction
of velocity. '' kc ''*KC20''

g TME0

```
do TME0_principle_selection
do TME0_write_eqn
do TME0_solve_equation
do TME0_discuss_qualitative
do TME0_discuss_relevance rand sem justify
```

g TME0_principle_selection

```
do TME0_principle_selection1
do TME0_principle_selection2
say ''[cont]In this problem, the system only has one object, the
    rock. Therefore, there are no potential energies involved,
    neither gravitational potential energy nor spring potential
    energy. '' kc ''*KC21 *KC22''
do TME0_discuss_truth rand
```

g TME0_principle_selection1 sem elicit

```
say ''First, let's find out TME0. Which principle will help you
    find TME0? ''
if ''Definition of total mechanical energy.'' true kc ''+KC24''
otherwise say ''[cont]let's use the definition of total
    mechanical energy.'' kc ''-KC24 *KC24''
```

```

g TME0_principle_selection1 sem tell
  say ''[cont]Let's use the definition of total mechanical energy
  to find TME0.'' kc ''*KC24''

g TME0_principle_selection2 sem elicit
  say ''[ask 'form'] What is the generic equation for the
  definition of total mechanical energy? TME=''
  if ''{form_TME}'' true kc ''+KC24q''
  otherwise say ''[cont]The generic equation for the definition of
  total mechanical energy is TME={form_TME}'' kc ''-KC24q *KC24q''

g TME0_principle_selection2 sem tell
  say ''[cont]The generic equation for the definition of total
  mechanical energy is TME={form_TME}.'' kc ''*KC24q''

g TME0_write_eqn
  say ''[ask 'eqn'] Please write the equation for the definition
  of total mechanical energy for the rock system at T0.''
  if ''{eqn_TME0}'' true say ''[add 'eqn_TME0']'' kc ''+KC24q''
  if ''equation with substituted values'' do bo_TME0_sub_write_eqn
  otherwise do bo_TME0_write_eqn kc ''-KC24q''

  g bo_TME0_sub_write_eqn
  say ''[cont]You substituted problem values into the equation.
  Please use the variable labels instead.''
  do TME0_write_eqn

  g bo_TME0_write_eqn
  say ''[add 'eqn_TME0']''
  say ''[cont]I have written the system's total mechanical energy

```

equation at T0 for you: {eqn_TME0}'' kc ''*KC24q''

g TME0_solve_equation sem elicit

say ''[ask 'num_units']It is easy to infer the value of TME0. Please calculate the value of TME0.''

if ''{TME0.value}'' true say ''[add 'TME0']'' kc ''+KC23''

if ''correct value, wrong/no unit'' say ''[cont]You entered the correct value, but not the correct units. The correct units for total mechanical energy are J. '' kc ''-KC23 *KC23''

otherwise do bo_TME0_solve_equation

say ''[add 'TME0']''

g bo_TME0_solve_equation

say ''[cont]It is easy. TME0 and KE0 have the same value: {TME0.value}'' kc ''*KC23''

g TME0_solve_equation sem tell

say ''[add 'TME0']''

say ''[cont]From {eqn_TME0}, it is easy to infer the value of TME0. TME0 and KE0 have the same value: {TME0.value}''

kc ''*KC23''

g TME0_discuss_qualitative

do TME0_discuss_qualitative_1

do TME0_discuss_qualitative_2 sem justify

g TME0_discuss_qualitative_1 sem elicit

say ''[menu 'Yes, it is true|No, it is not true|I do not know|No
t enough information']Equation {eqn_TME0.value} tells us that
the rock system's total mechanical energy is equal to the rock
's kinetic energy at T0. Is it true *for any given time*?''
if ''Yes, it is true.'' true kc ''+KC24''
otherwise say ''[cont]In fact ...it is true for any time point
in this problem. '' kc ''-KC24 *KC24''

g TME0_discuss_qualitative_1 sem tell

say ''[cont]Equation {eqn_TME0.value} tells us that the system's
total mechanical energy is equal to the rock's kinetic energy
at T0. In fact, this is true for any given time point in
this problem.'' kc ''*KC24''

g TME0_discuss_qualitative_2 sem elicit

say ''Why?''

if ''Because as long as the rock is the only object in the
system, there are no potential energies involved. Thus the total
mechanical energy equals the rock's kinetic energy at any given
time point. '' true kc ''+KC24 +KC21 +KC22''

otherwise say ''[cont]This is because as long as the rock is the
only object in the system, there are no potential energies
involved, neither gravitational potential energy nor spring
potential energy. Thus the rock system's total mechanical energy
equals the rock's kinetic energy at any given time point. '' kc
''-KC24 -KC21 -KC22 *KC24 *KC21 *KC22''

g TME0_discuss_qualitative_2 sem tell

say ''[cont]This is because as long as the rock is the only
object in the system, there are no potential energy involved,

```
neither gravitational potential energy nor spring potential energy
. Thus TME = KE at any given time point. ''
kc ''*KC24 *KC21 *KC22''
```

```
g TME0_discuss_truth
do TME0_discuss_truth1 sem justify

g TME0_discuss_truth1 sem elicit
say ''Why are there no potential energies involved in this
problem?'' rand
if ''Because the rock is the only object in the system, there
are no potential energies involved. '' true kc ''+KC21''
otherwise say ''[cont]Recall that potential energy is associated
with the relative positions of two objects. Because the rock
is the only object in the system, there are no potential
energies involved. '' kc ''-KC21 *KC21''

g TME0_discuss_truth1 sem tell
say ''[cont]Next, I want to explain why there are no potential
energies involved in this problem.''
say ''[cont]Recall that potential energy is associated with the
relative positions of two objects. There is only one object
in this problem, the rock, and thus potential energy is not
involved in this problem. '' kc ''*KC21''

g TME0_discuss_relevance sem elicit
say ''[cont]Let's go back to continue our problem solving.''
```

```

say ''Thinking from the perspective of the whole solution plan,
why do we need to calculate the system's total mechanical
energy at T0?'' rand
if ''We can calculate the net work done on the rock from T0 to
T1. '' true kc ''+KC28''
otherwise say ''[cont]Because if we already know TME0 and TME1,
we can calculate the net work on the rock from T0 to T1.'' kc ''
-KC28 *KC28''

```

```

g TME0_discuss_relevance sem tell
say ''[cont]Let's go back to continue our problem solving.''
say ''[cont]Thinking from the perspective of the whole solution
plan, the reason that we need to calculate the system's total
mechanical energy at T0 is so that we can calculate the net
work done on the rock from T0 to T1. '' kc ''*KC28''

```

```

g TME1
do TME1_principle_selection
do TME1_write_eqn
do TME1_solve_equation
do TME1_discuss_truth rand
do TME1_discuss_qualitative rand

```

```

g TME1_principle_selection sem elicit
say ''[cont]Since the total mechanical energy is equal to
kinetic energy at any given time as long as the rock is the only
object in a system, it is simple to infer the value for TME1.''
kc ''*KC24''
say ''Which principle will help you find TME1? ''

```

```

if ''Definition of total mechanical energy. '' true kc ''+KC24''
otherwise say ''[cont]Let's use the definition of total
mechanical energy again.'' kc ''-KC24 *KC24''

g TME1_principle_selection sem tell
say ''[cont]Since the total mechanical energy is equal to
kinetic energy at any given time as long as the rock is the only
object in a system, let's use the definition of total mechanical
energy again to get the value for TME1.'' kc ''*KC24''

g TME1_write_eqn
say ''[ask 'eqn'] Please write the appropriate equation for the
definition of total mechanical energy of the system at T1.''
if ''{eqn_TME1}'' true say ''[add 'eqn_TME1']'' kc ''+KC24q''
if ''equation with substituted values'' do bo_TME1_sub_write_eqn
otherwise do bo_TME1_write_eqn kc ''-KC24q''

g bo_TME1_sub_write_eqn
say ''[cont]You substituted problem values into the equation.
Please use the variable labels instead.''
do TME1_write_eqn

g bo_TME1_write_eqn
say ''[add 'eqn_TME1']''
say ''[cont]Let me just write the system's total mechanical
energy equation at T1 for you: {eqn_TME1.value}.'' kc ''*KC24q''

g TME1_solve_equation sem elicit
say ''[ask 'num_units']What is the value of TME1?''
if ''{TME1.value}'' true say ''[add 'TME1']'' kc ''+KC23''

```

```

if ''correct value, wrong/no unit'' say ''[cont]You entered the
  correct value, but not the correct units. The correct units for
  total mechanical energy are J. '' kc ''-KC23 *KC23''
otherwise do bo_TME1_solve_equation
say ''[add 'TME1']''

g bo_TME1_solve_equation
say ''[cont]It is easy. TME1 and KE1 have the same value: {TME1.
  value}'' kc ''*KC23''

g TME1_solve_equation sem tell
say ''[add 'TME1']''
say ''[cont]We can easily get the value of TME1. TME1 and KE1
  have the same value: {TME1.value} '' kc ''*KC23''

g TME1_discuss_truth
do TME1_discuss_truth1_1
do TME1_discuss_truth1_2 sem justify
do TME1_discuss_truth2 sem justify

g TME1_discuss_truth1_1 sem elicit
say ''[cont]Ok, before going on to the next step, let me ask you
  a few questions about the definition of total mechanical energy.''
say ''[menu 'Yes, it can be negative|No, it cannot be negative|I
  do not know|Not enough information']In this problem, can the
  system's total mechanical energy at any given time be negative? ''
if ''No, it cannot be negative.'' true
otherwise say ''[cont]Actually ... the system's total mechanical

```

energy at any given time cannot be negative. ''

g TME1_discuss_truth1_1 sem tell

say ''[cont]Okay, before going on to the next step, let me tell you a few more things about the definition of total mechanical energy.''

say ''[cont]In this problem, the rock system's total mechanical energy at any given time cannot be negative. '' kc ''*KC24''

g TME1_discuss_truth1_2 sem elicit

say ''Please explain why.''

if ''Because TME = KE at any given time point, and the rock's kinetic energy cannot be negative.''' true kc ''+KC24 +KC20''

otherwise say ''[cont]This is because the rock system's total mechanical energy equals the rock's kinetic energy at any given time point, and the rock's kinetic energy cannot be negative. '' kc ''-KC24 -KC20 *KC24 *KC20''

g TME1_discuss_truth1_2 sem tell

say ''[cont]This is because the rock system's total mechanical energy equals the rock's kinetic energy at any given time point in this problem, therefore, the rock system's total mechanical energy at any given time cannot be negative. '' kc ''*KC24 *KC20''

g TME1_discuss_truth2 sem elicit

say ''Explain why the rock's kinetic energy at any time point cannot be negative?''

if ''Because KE={form_KE.value}, and neither the rock's mass

```

    nor the  $v^2$  can be negative, then kinetic energy cannot be
    negative. '' true kc ''+KC20q +KC20''
otherwise say ''[cont]Because  $KE=\{form\_KE.value\}$ . Neither the ro
    ck's mass nor  $v^2$  can be negative, then kinetic energy cannot
    be negative. '' kc ''-KC20 *KC20q *KC20''

g TME1_discuss_truth2 sem tell
say ''[cont]The reason that the rock's kinetic energy cannot be
    negative is because  $KE=\{form\_KE.value\}$ . Neither the rock's
    mass nor  $v^2$  can be negative.'' kc ''*KC20q *KC20''

g TME1_discuss_qualitative
    do TME1_discuss_qualitative_1
    do TME1_discuss_qualitative_2 sem justify

g TME1_discuss_qualitative_1 sem elicit
say ''[menu 'Yes, we can|No, we cannot|I do not know|Not enough
information']Given the system's total mechanical energy is
    equal to rock's kinetic energy at any given time point
    in this problem, can we say anything about the ROCK's
    total mechanical energy, instead of the system's total
    mechanical energy?''
    if ''No, we cannot.'' true kc ''+KC24''
otherwise say ''[cont]Actually... No, we cannot.''
    kc ''-KC24 *KC24''

g TME1_discuss_qualitative_1 sem tell
say ''[cont]Although the system's total mechanical energy is equal
    to the rock's kinetic energy at any given time point in this
    problem, we cannot refer to the rock's total mechanical energy.

```

```

''kc ''*KC24''

g TME1_discuss_qualitative_2 sem elicit
say ''Why?''
    if ''Because the TME always refers to a system. In the problem,
        we only have one object (the rock) in the system. But it
        still only refers to the system of the rock, not the
            object rock.'' true kc ''+KC24''
otherwise say ''[cont]This is because the total mechanical energy
always refers to a system. In the problem, we only have one
object in the system. But total mechanical energy must
still only refers to the system of the rock, not the
object rock.'' kc ''-KC24 *KC24''

g TME1_discuss_qualitative_2 sem tell
say ''[cont]This is because the total mechanical energy always
refers to a system. Here we only have one object in the
system, but total mechanical energy must still only
refers to the system of the rock, not the object
rock.'' kc ''*KC24''

g CME01
do CME01_principle_selection
do CME01_argument_selection
do CME01_write_eqn
do CME01_solve_equation
say ''[cont]We have finished solving the whole problem. Next let's do
a little discussion.''

```

```

g CME01_principle_selection
  do CME01_principle_selection2
  do CME01_principle_selection3

g CME01_principle_selection2 sem elicit
  say ''Which principle will help you calculate the work done on
  the rock from T0 to T1? Please provide the name of the
  principle, not an equation. ''
  if ''conservation of mechanical energy for non-isolated systems''
    true kc ''+KC28''
  otherwise do temp-CME01_principle_selection2 kc ''-KC28''

  g temp-CME01_principle_selection2 sem tell
    say ''[cont]As we have mentioned before, there were two
    principles from the textbook which apply to net
    work: one is the definition of net work and
    the other is conservation of mechanical energy
    for non-isolated systems.'' kc ''*KC28 *KC18''
  say ''[cont]The definition of the net work cannot be
    applied to this problem because there is no
    information given in the problem statement about
    the forces on the rock. Therefore, we will apply the
    remaining principle, namely, conservation of
    mechanical energy for non-isolated systems.
    We will start with the generic equation.''
    kc ''*KC18 *KC28''

g CME01_principle_selection2 sem tell

```

```
say ''[cont]As we have mentioned before, there were two principles from the textbook which apply to net work: one is the definition of net work and the other is conservation of mechanical energy for non-isolated systems.'' kc ''*KC28 *KC18''
```

```
say ''[cont]The definition of the net work cannot be applied to this problem because there is no information given in the problem statement about the forces on the rock. Therefore, we will apply the remaining principle, namely, conservation of mechanical energy for non-isolated systems. We will start with the generic equation.'' kc ''*KC18 *KC28''
```

```
g CME01_principle_selection3 sem elicit
```

```
say ''[ask 'form'] What is the generic equation for conservation of mechanical energy for non-isolated systems?''
```

```
if ''{form_Wnet01}'' true kc ''+KC28q''
```

```
otherwise say ''[cont]The generic equation for conservation of mechanical energy for non-isolated systems: {form_Wnet01}'' kc ''-KC28q *KC28q''
```

```
g CME01_principle_selection3 sem tell
```

```
say ''[cont]The generic equation for conservation of mechanical energy for non-isolated systems is: {form_Wnet01}'' kc ''*KC28q''
```

```
g CME01_argument_selection sem elicit
```

```
say ''What time points and bodies should we use when applying the principle of conservation of mechanical energy for non-isolated systems?'' answer ''T0 and T1'' answer ''system''
```

```
if ''T0 and T1'' true do-nomatch CME01_argument_miss_T kc ''+KC28''
```

```
if ''system'' true do-nomatch CME01_argument_miss_object kc ''+KC28''
```

```

    if ''rock'' false do CME01_argument_rock_system
    otherwise do CME01_argument_miss_both kc ''-KC28''

g CME01_argument_miss_both_time
  say ''[cont]To apply conservation of mechanical energy for
  non-isolated systems, we need two time points.''
  say ''What time points should we use when applying
  conservation of mechanical energy for non-isolated
systems?''
  if ''T0 and T1'' true kc ''+KC28''
  otherwise say ''[cont]We should apply conservation of
  mechanical energy for non-isolated systems at TWO
time points: T0 and T1.''
  kc ''-KC28 *KC28''

g CME01_argument_miss_T
  say ''[cont]We should apply conservation of mechanical
  energy for non-isolated systems at TWO time points:
T0 and T1.'' kc ''*KC28''

g CME01_argument_miss_object
  say ''What objects should you apply the conservation of
  mechanical energy for non-isolated systems to?''
  opt sem already-rock-system
  if ''system'' true kc ''+KC28''
  if ''rock'' false do CME01_argument_rock_system
  otherwise say ''[cont]We have already defined the system
  to be the rock. Therefore, we should apply
conservation of mechanical energy to the
rock system.'' kc ''-KC28 *KC28''

```

```

g CME01_argument_rock_system
say ''Does the rock refer to the object or the system?''
  sem already-rock-system
if ''system '' true kc ''+KC28''
otherwise say ''[cont]We always use the term the total
  mechanical energy in relation to a system. In this
problem, the rock is the system. Therefore, we
should apply conservation of mechanical
  energy to the rock system. '' kc ''-KC28 *KC28''

g CME01_argument_miss_both
do CME01_argument_miss_object
do CME01_argument_miss_both_time

g CME01_argument_selection sem tell
say ''[cont]We should apply conservation of mechanical energy
  for non-isolated systems for the rock system on the time points:
  T0 and T1.'' kc ''*KC28''

g CME01_write_eqn sem elicit
say ''[ask 'eqn']Please write the equation for conservation of
  mechanical energy for non-isolated systems for the rock system
  on the time points: T0 and T1.''
if ''{eqn_CME01}'' true say ''[add 'eqn_CME01']'' kc ''+KC28q''
if ''equation with substituted values'' do bo_CME01_sub_write_eqn
otherwise do bo_CME01_write_eqn kc ''-KC28q''

g bo_CME01_sub_write_eqn

```

```

say ''[cont]You substituted problem values into the equation.
    Please use the variable labels instead.'''
do CME01_write_eqn

g bo_CME01_write_eqn
say ''[add 'eqn_CME01']''
say ''[cont]Let me just write the equation for conservation of
mechanical energy for the rock system between T0 and T1 for you:
{eqn_CME01}.'' kc ''*KC28q''

g CME01_write_eqn sem tell
say ''[add 'eqn_CME01']''
say ''[cont]Let me just write the equation for conservation of
mechanical energy for the rock system between T0 and T1 for you:
{eqn_CME01}.'' kc ''*KC28q''

g CME01_solve_equation sem elicit
say ''[cont]Since we already know the values for TME0 and TME1,
we can calculate the value of Wnet01 easily.'''
say ''[ask 'num_units']Please calculate the value of Wnet01.'''
if ''{Wnet01.value}'' true say ''[add 'Wnet01']'' kc ''+KC15''
if ''correct value, wrong/no unit'' say ''[cont]You entered the
correct value, but not the correct units. The correct units for
net works are J. '' kc ''-KC15 *KC15''
otherwise do bo_CME01_solve_equation
say ''[add 'Wnet01']''

g bo_CME01_solve_equation
say ''[cont]Based on: {eqn_CME01}, TME1={TME1.value}, and TME0= {
TME0.value}, we get: {TME0.value}+Wnet01={TME1.value}. That is:

```

Wnet01={TME1.value}-{TME0.value}. I calculate the value for
Wnet01 to be {Wnet01.value}. '' kc ''*KC15''

g CME01_solve_equation sem tell

say ''[cont]Since we already know the values for TME0 and TME1,
we can calculate the value of Wnet01 easily.''

say ''[add 'Wnet01']''

say ''[cont]Based on: {eqn_CME01}, TME1={TME1.value}, and TME0=
{TME0.value}, we get: {TME0.value}+Wnet01={TME1.value}. That
is: Wnet01={TME1.value}-{TME0.value}. I calculate the value for
Wnet01 to be {Wnet01.value}. '' kc ''*KC15''

g discuss_problem

do extra_step

say ''[cont]The reason that we calculate v1 in this problem is to
practice the principle of the definition of kinetic energy. ''

do PE_concept_justification

say ''[cont]Next, we will change the place where the rock is and
discuss whether the values of some variables would also change.''

do v1-ppd-discussion

do Wnet01-ppd-discussion

do inc_dec

say ''[cont]Next, we will change the mass of the rock and discuss
whether the values of some variables would also change.''

do inc_mass

do work_concept sem justify

do compare_path

g extra_step sem elicit

```

say ''[menu 'KE0|v1|TME1|TME0|Wnet01|All are necessary|I do not know
']While solving the problem, we calculated the variables in the
order of KE0, v1, TME0, TME1, and Wnet01. Are any of these
variable(s) not necessary to solve the problem?''
if ''v1'' true
otherwise say ''[cont]Actually ... v1 was not necessary for solving
the problem. ''

g extra_step sem tell
say ''[cont]While solving the problem, we calculated the variables
in the order of KE0, v1, TME0, TME1, and Wnet01. However, v1 was
not necessary for solving the problem. ''

g PE_concept_justification
do PE_concept_justification_1
do PE_concept_justification_2 rand sem justify

g PE_concept_justification_1 sem elicit
say ''[menu 'Yes, it is possible|No, it is not possible|I do not
know|Not enough information']In this problem, we have selected the
rock as the system. Is it possible to define potential energy for
the rock system?''
if ''No, it is not possible'' true kc ''+KC21''
otherwise say ''[cont]Actually ... it is not possible to define
potential energy for the rock system.'' kc ''-KC21 *KC21''

g PE_concept_justification_1 sem tell

```

```
say ''[cont]In this problem, we have selected the rock as the system
and we *cannot* define a potential energy for the rock system.
'' kc ''*KC21''
```

```
g PE_concept_justification_2 sem elicit
```

```
say ''Why not?''
```

```
if ''Potential energy can be defined only for a system which contains
two or more objects but here we only have only one object as the
system (the rock).'' true kc ''+KC21''
```

```
otherwise say ''[cont]This is because potential energy is energy
associated with the configuration of two or more objects. Thus
potential energy can be defined only for a system which contains
two or more objects. Here we only have the rock to select as the
system and thus we cannot define a potential energy for the rock.'''
kc ''-KC21 *KC21''
```

```
g PE_concept_justification_2 sem tell
```

```
say ''[cont]This is because potential energy is an energy associated
with the configuration of two or more objects. Thus potential
energy can be defined only for a system which contains two
or more objects.''' kc ''*KC21''
```

```
say ''[cont]In this problem, we only have the rock to select as the
system. Therefore, we cannot define a potential energy for the
rock. ''
```

```
g v1-ppd-discussion
```

```
do v1-ppd-discussion-1
```

```
do v1-ppd-discussion-2 rand sem justify
```

g v1-ppd-discussion-1 sem elicit

say ''[menu 'increase|decrease|same|not enough information|I do not know']Suppose the 0.6 kg rock is put on the Earth and the rock's kinetic energy at point B is still 7.5 J, what will happen to the rock's magnitude of velocity at the point B?''

if ''same'' true kc ''+KC20''

if ''{v1.mag}'' true kc ''+KC20 +KC31''

otherwise say ''[cont]Actually ... the rock's magnitude of velocity should be same, still be {v1.mag}. '' kc ''-KC20 *KC20 *KC31''

g v1-ppd-discussion-1 sem tell

say ''[cont]Suppose the rock is put on the Earth and the rock's kinetic energy at point B is still 7.5 J, then the rock's magnitude of velocity at the point B is still {v1.mag}.'' kc ''*KC20 *KC31''

g v1-ppd-discussion-2 sem elicit

say ''Why?''

if ''Because if the rock's KE at point B and the rock's mass stay the same, the rock's magnitude of velocity at the point B is the same.'' true kc ''+KC20''

otherwise say ''[cont]Because if the rock's kinetic energy at point B and the rock's mass stay the same, then according to $KE = \frac{1}{2}mv^2$ the rock's magnitude of velocity at the point B is the same.'' kc ''-KC20 *KC20q *KC20''

g v1-ppd-discussion-2 sem tell

say ''[cont]Because the kinetic energy stays the same and the rock's mass stays the same, then according to $KE = \frac{1}{2}mv^2$ the rock's magnitude of velocity at the point B is same.'' kc ''*KC20q *KC20''

```

g Wnet01-ppd-discussion
do Wnet01-ppd-discussion-1
do Wnet01-ppd-discussion-2 rand sem justify

g Wnet01-ppd-discussion-1 sem elicit
say ''[menu 'increase|decrease|same|not enough information|I do not
know.')]Suppose the rock is on the Earth and everything in the
problem statement is kept the same. What will happen to net work
done on the rock if we still choose the rock alone as the system?''

if ''same'' true kc ''+KC28 +KC24''
otherwise say ''[cont]Actually ... It should be the same:{Wnet01.va
lue}.'' kc ''-KC28 -KC24 *KC28 *KC24 *KC15''

g Wnet01-ppd-discussion-2 sem elicit
say ''Why?'' rand
if ''Since the rock is the only object in the system. The TME = KE
any given time. The rock's KE0 and KE1 are still the same as
those when the rock is in space. Therefore, the TME1 and TME2 are the
same as those when the rock is in space. Therefore, the Wnet01 i
s still the same as when the rock is in space. '' true kc ''+KC28
+KC24q''
otherwise do tepm-Wnet01-ppd-discussion kc ''-KC28 -KC24''

g tepm-Wnet01-ppd-discussion
say ''[cont]Let me explain it to you step by step. ''
say ''[cont]The rock is still the only object in the system and thus
the systems' total mechanical energy is still equal to its
kinetic energy at any given time. '' kc ''*KC24''

```

say ''[cont]The rock's kinetic energy at T0 and T1 are still the same as those when the rock is in space because m , v_0 , and KE_1 is kept the same. Therefore, the system's TME1 and TME2 are the same as those when the rock is in space. '' kc ''*KC20''

say ''[cont]As a result, the net work done on the rock from T0 to T1 is still the same as when the rock is in space. ''

kc ''*KC28''

g Wnet01-ppd-discussion-1 sem tell

say ''[cont]Suppose the rock is on the Earth and everything in the problem statement is kept the same. If we still choose the rock alone as the system, then the net work done on the rock from T0 to T1 is still $\{W_{net01.value}\}$. '' kc ''*KC28 *KC24 *KC15''

g Wnet01-ppd-discussion-2 sem tell

say ''[cont]Let me explain it to you step by step. ''

say ''[cont]Since the rock is still the only object in the system and thus the systems' total mechanical energy is still equal to its kinetic energy at any given time. '' kc ''*KC24''

say ''[cont]The rock's kinetic energy at T0 and T1 are still the same as those when the rock is in space because m , v_0 , and KE_1 is kept the same. Therefore, the system's TME1 and TME2 are the same as those when the rock is in space. '' kc ''*KC20''

say ''[cont]As a result, the net work done on the rock from T0 to T1 is still the same as when the rock is in space. '' kc ''*KC28''

g inc_dec

```

say ''[ask 'num_units']Let's go back to the situation where rock is
  in space. Look at the variable window, tell me what is the
  magnitude of v0?''
if ''{v0.mag}'' true kc ''+KC31''
if ''correct value, wrong/no unit'' say ''[cont]You entered the
  correct value, but not the correct units. The correct units for
  velocity are m/s. '' kc ''-KC31 *KC31''
otherwise say ''[cont]We can also see from the variable window that
  v0 has a value of 2.0 m/s.'' kc ''*KC31''
say ''[ask 'num_units']What is the magnitude of v1?''
if ''{v1.mag}'' true kc ''+KC31''
if ''correct value, wrong/no unit'' say ''[cont]You entered the
  correct value, but not the correct units. The correct units for
  velocity are m/s. '' kc ''-KC31 *KC31''
otherwise say ''[cont]We can also see from the variable window that
  v1 has a value of 5.0 m/s.'' kc ''*KC31''
do v1_discuss_qualitative3
say ''[ask 'num_units']Look at the variable window, tell me what is
  the value of the rock's kinetic energy at T0?''
if ''{KE0.value}'' true kc ''+KC23''
if ''correct value, wrong/no unit'' say ''[cont]You entered the
  correct value, but not the correct units. The correct units for
  kinetic energy are J. '' kc ''-KC23 *KC23''
otherwise say ''[cont]We can also see from the variable window that
  KE0 has a value of {KE0.value}.'' kc ''*KC23''
say ''[ask 'num_units']What is the value of KE1?''
if ''{KE1.value}'' true kc ''+KC23''
if ''correct value, wrong/no unit'' say ''[cont]You entered the
  correct value, but not the correct units. The correct units for
  kinetic energy are J. '' kc ''-KC23 *KC23''

```

```

otherwise say ''[cont]We can also see from the variable window that
    KE1 has a value of {KE1.value}.'' kc ''*KC23''
do compare_A_B
do compare_C_A
do compare_C_B

g compare_A_B
do compare_A_B_1
do compare_A_B_2 sem justify

g compare_A_B_1 sem elicit
say ''[menu 'increase|decrease|same|not enough information|I do not
    know.']]Compared with point A, what can you say about the kinetic
    energy of the rock at point B?''
if ''Increases'' true kc ''+KC20''
otherwise say ''[cont]Actually ... It is increased.'' kc ''-KC20 *KC2
    0''

g compare_A_B_1 sem tell
say ''[cont]Compared with point A, the kinetic energy of the rock
    at point B is increased.'' kc ''*KC20''

g compare_A_B_2 sem elicit
say ''Please try to explain why.''
if ''Because {KE1.value} is bigger than {KE0.value}, therefore, it
    is increased. '' true kc ''+KC20 +KC23''
otherwise say ''[cont]This is a bit of a trick question. The problem
    doesn't say why the kinetic energy increase, but instead just
    gives initial conditions (v0=2 m/s) and final conditions (KE1= 7.5
    J) that imply that the kinetic energy has increased because {KE1

```

.value} is bigger than {KE0.value}.'' kc ''*KC23 *KC31''

g compare_A_B_2 sem tell

say ''[cont]The problem doesn't say why the kinetic energy increase
but we can infer it from the fact that KE1= {KE1.value} is
bigger than the value of KE0: {KE0.value}.'' kc ''*KC23''

g v1_discuss_qualitative3 sem elicit

say ''[menu 'increase|decrease|same|not enough information|I do
not know']What can we infer about that the rock's magnitude of
velocity during the ENTIRE period from T0 to T1 given that
the magnitude of v1 (5.0 m/s) is bigger than the magnitude of v0
(2.0 m/s)?'' kc ''*KC31''

if ''not enough information.'' true kc ''+KC20''

otherwise say ''[cont]We do not have enough information to
deduce how the rock's magnitude of velocity changes during the
process from T0 to T1. '' kc ''-KC20 *KC20''

g v1_discuss_qualitative3 sem tell

say ''[cont]Although the magnitude of v1 (5.0 m/s) is bigger
than the magnitude of v0 (2.0 m/s), we *cannot* infer that the
rock's magnitude of velocity is increasing during the ENTIRE
period from T0 to T1. '' kc ''*KC31 *KC20''

g compare_C_A

do compare_C_A_1

do compare_C_A_2 sem justify

g compare_C_A_1 sem elicit

say ''[menu 'increase|decrease|same|not enough information|I do not

know.']Suppose there is point C on the path of the rock between A to B. Compared with point A, what can you say about the kinetic energy of the rock at point C?''

if ''Not enough information.'' true kc ''+KC20''

otherwise say ''[cont]Actually ... We do not have enough information to say about the kinetic energy of the rock at point C. '' kc ''-KC20 *KC20''

g compare_C_A_1 sem tell

say ''[cont]Suppose there is point C on the path of the rock between A to B. We do not have enough information to say anything about the kinetic energy of the rock at point C compared with point A . '' kc ''*KC20''

g compare_C_A_2 sem elicit

say ''Please explain why.''

if ''We do not know whether the rock's magnitude of velocity at C is bigger than its magnitude of velocity at A or not.'' true kc ''+KC20''

otherwise say ''[cont]We do not know whether the rock's magnitude of velocity at C is bigger than its magnitude of velocity at A or not. It is possible that the rock continuously increases its magnitude of velocity from A to C then to B. It is also possible its magnitude of velocity decreases from A to C and then increases magnitude of velocity from C to B. It is even possible that it does not change its magnitude of velocity from A to C and then increase its magnitude of velocity from C to B. Therefore, we do not have enough information to say about the

relationship between the kinetic energy of the rock at point C and at point A.''' kc ''-KC20 *KC20''

g compare_C_A_2 sem tell

say ''[cont]This is because we do not know whether the rock's magnitude of velocity at C is bigger than its magnitude of velocity at A or not. It is possible that the rock continuously increases its magnitude of velocity from A to C then to B. It is also possible its magnitude of velocity decreases from A to C and then increases from C to B. It is even possible that it does not change its magnitude of velocity from A to C and then increases its magnitude of velocity from C to B.''' kc ''*KC20''

g compare_C_B

do compare_C_B_1

do compare_C_B_2 sem justify

g compare_C_B_1 sem elicit

say ''[menu 'increase|decrease|same|not enough information|I do not know.']*Compared with point B, What can you say about the kinetic energy of the rock at point C?''

if ''Not enough information.''' true kc ''+KC20''

otherwise say ''[cont]Actually ... We do not have enough information to say about the kinetic energy of the rock at point C compared with point B either. '' kc ''-KC20 *KC20''

g compare_C_B_1 sem tell

say ''[cont]Similarly, we do not have enough information to say any thing about the kinetic energy of the rock at point C compared with point B either.'' kc ''*KC20''

g compare_C_B_2 sem elicit

say ''Please explain why.''

if ''We do not have enough information about the relationship between the rock's magnitude of velocity at point C and at B.''

kc ''+KC20''

otherwise say ''[cont]The answer is similar to reasons that we do not have enough information about the relationship between the rock's magnitude of velocity at point A and point C, it is because we do not know whether the magnitude of the rock's magnitude of velocity at C is bigger than its magnitude of velocity at B or not. That is, we do not have enough information about the relationship between the rock's magnitude of velocity at point C and at B.''

kc ''-KC20 *KC20''

g compare_C_B_2 sem tell

say ''[cont]Because we do not know whether the rock's magnitude of velocity at C is bigger than its magnitude of velocity at B or not. That is, we do not have enough information about the relationship between the rock's magnitude of velocity at point C and at B. ''

kc ''*KC20''

g work_concept

```
do work_concept_whether
do work_concept_justification rand sem justify
```

```
g work_concept_whether sem elicit
say ''[menu 'Yes, we can|No, we cannot|not enough information|
  I do not know.']*If we already know that the work done on the rock
  by force F is 9 J and we also know that the rock is moving in the
  same direction as F for 3 m in a straight line, can you infer the
  magnitude of the force F?'' kc ''*KC15''
if ''No, we cannot.'' true
otherwise say ''[cont]No, we cannot infer anything about the
  magnitude of the force F. ''
```

```
g work_concept_whether sem tell
say ''[cont]If we already know that the work done on the rock by
  force F is 9 J and we also know that the rock is moving in the
  same direction as F for 3 m in a straight line,
  we *cannot* infer the magnitude of the force F.''' kc ''*KC14 *KC15''
```

```
g work_concept_justification sem elicit
say ''Why not?''
if ''Because we do not know whether F is a constant force or not.'''
  true kc ''+KC14''
otherwise say ''[cont]It is because we do not know whether F is a
  constant force or not. If F is a constant force, we can apply
   $W=F*d$  to get that  $F = W/d = (9 \text{ J})/(3 \text{ m}) = 3 \text{ N}$ ; however, in most
  cases, F is not constant force. Therefore, we *cannot* apply
   $W=F*d$  to get the magnitude of the F.'''
kc ''-KC14 *KC14q *KC14 *KC15''
```

```

g work_concept_justification sem tell
  say ''[cont]It is because we do not know whether F is a constant
    force or not. Only if F is a constant force, we can apply  $W=F*d$ 
    to get that  $F = W/d = (9 \text{ J})/(3 \text{ m}) = 3 \text{ N}$ ; however, in most cases,
    F is not constant force. Therefore, we cannot apply  $W=F*d$  to
    get the magnitude of the F.'' kc ''*KC14 *KC14q *KC15''

g inc_mass
  do inc_mass1_1
  do inc_mass1_2 rand sem justify
  do inc_mass2_1
  do inc_mass2_2 rand sem justify
  do inc_mass3_1
  do inc_mass3_2 rand sem justify

g inc_mass1_1 sem elicit
  say ''[menu 'increase|decrease|same|not enough information|
    I do not know.']*Suppose the rock's mass increased from 0.6 kg to
    1 kg and the rock's kinetic energy at point B was still 7.5 J,
    what would happen to the rock's magnitude of velocity at point B?''
  kc ''*KC23''
  if ''decrease'' true kc ''+KC20''
  otherwise say ''[cont]Actually ... the rock's magnitude of velocity
    decreases at point B. '' kc ''-KC20 *KC20''

g inc_mass1_1 sem tell
  say ''[cont]Suppose that the rock's mass increased from 0.6 kg to 1

```

kg and the rock's kinetic energy at point B was still 7.5 J,
the rock's magnitude of velocity at point B would decrease. '''
kc '''*K C20 *KC23''

g inc_mass1_2 sem elicit

say '''Why?''

if '''It decreases because if the rock's KE at point B stays same,
{KE1.value} and the rock's mass increased from 0.6 kg to 1 kg.
According to KE={form_KE.value}, the magnitude of velocity would
have to decrease.''' true kc '''+KC20q +KC20 +KC23''

otherwise say '''[cont]It decreases because we are assuming that the
rock's kinetic energy at point B stays same, {KE1.value} but the
rock's mass increased from 0.6 kg to 1 kg. According to KE=
{form_KE.value}, the magnitude of velocity would have to
decrease.''' kc '''-KC20 *KC20q *KC20 *KC23''

g inc_mass1_2 sem tell

say '''[cont]It is because we are assuming that the rock's kinetic
energy at point B stays same, {KE1.value}, and if the rock's
mass increases, according to KE={form_KE}, then the magnitude
of rock's magnitude of velocity has to decrease.'''

kc '''*KC20 *KC20q *KC23''

g inc_mass2_1 sem elicit

say '''[menu 'increase|decrease|same|not enough information|

I do not know.']. Suppose the rock's mass were 1 kg instead
of 0.6 kg and yet the rock's magnitude of velocity at point
A was still 2.0 m/s. What will happen to the rock's kinetic
energy at point A?'' kc '''*KC31''

```

if ''Increase'' true kc ''+KC20''
otherwise say ''[cont]Actually ... the rock's kinetic energy at
point A increases.'' kc ''-KC20 *KC20''

g inc_mass2_1 sem tell
say ''[cont]Suppose the rock still has a velocity of magnitude of
2.0 m/s at point A but the rock's mass increased from 0.6 kg
to 1 kg, then the rock's kinetic energy at point A should
increase.'' k c ''*KC31 *KC20''

g inc_mass2_2 sem elicit
say ''Why?'' rand
if ''Because the magnitude of velocity stays the same and the
mass increases, therefore the kinetic energy at point A
increases.'' true kc ''+KC20''
otherwise say ''[cont]Because the magnitude of rock magnitude
of velocity is still 2.0 m/s at point A and the mass of
the rock increases, according to  $KE=\{form\_KE.value\}$  the
kinetic energy of the rock at point A increases.''
kc ''-KC20 *KC20 *KC20q *KC31''

g inc_mass2_2 sem tell
say ''[cont]Because the rock's magnitufe of velocity is still
2.0 m /s at point A and the mass of the rock increases,
according to  $KE =\{form\_KE.value\}$  the kinetic energy of
the rock at point A increases.'' kc ''*KC20 *KC20q *KC31''

g inc_mass3_1 sem elicit
say ''[menu 'increase|decrease|same|not enough information|I do not
know.']*Finally, if the rock's mass were 1 kg instead of 0.6 kg

```

and the rock's magnitude of velocity at point A was still 2.0 m/s and the rock's kinetic energy at point B was still 7.5 J, what will happen to the net work done on the rock as it moves from A to B?" kc '*KC31 *KC23'

if 'Decrease' true kc '+KC28 +KC24 +KC20'

otherwise say '[cont]Actually ... the net work done on the rock as it moves from A to B will *decrease*.' kc '-KC28 -KC24 -KC20 *K C28 *KC24 *KC20'

g inc_mass3_2 sem elicit

say 'Why?'

if 'Because the rock's kinetic energy at point A increases and its kinetic energy at point B is same. Therefore, the net work done on the rock from A to B has to decrease.' true kc '+KC28 +KC24 +KC20'

otherwise do temp-inc_mass3 kc '-KC28 -KC24 -KC20'

g temp-inc_mass3

say '[cont]This is a tricky question. So I will explain it to you step by step.'

say '[cont]Given that the rock still has a velocity of magnitude 2.0 m/s at point A and a kinetic energy of 7.50

J at point B, the rock's kinetic energy at point A increases when the rock's mass increases from 0.6 kg to 1 kg. The new value is $0.5 \cdot 1 \text{ kg} \cdot (v_0)^2 = 2 \text{ J}$ while the old value is $0.5 \cdot 0.6 \text{ kg} \cdot (v_0)^2 = \{KE0.value\}$ ' kc '*KC20 *KC20q *KC23 *KC31'

say '[cont]In other words, the system's new TME1 is the same as the old TME1, which is $\{TME1.value\}$ but the system's

new TME0 is 2 J, which is bigger than the old TME0,
 which is {TME0.value}.’’
 kc ’’*KC24 *KC23’’
 say ’’[cont] According to {eqn_CME01}, we have $W_{net01} = TME1 -$
 TME0. When the rock’s mass increased from 0.6 kg to
 1 kg, the TME0 increased but TME1 remained the same,
 therefore, W_{net01} must decrease. That is, the net work
 done on the rock as it moves from A to B *decreases*
 as the rock’s mass increases from 0.6 kg to 1 kg.’’
 kc ’’*KC28q *KC28’’

g inc_mass3_1 sem tell

say ’’[cont]Finally, if the rock’s mass were 1 kg instead of 0.6 kg
 and the rock’s magnitude of velocity at point A was still
 2.0 m/ s and the rock’s kinetic energy at point B was still
 7.5 J, the net work done on the rock as it moves from A to B
 would decrease.’’
 kc ’’*KC28 *KC24 *KC20 *KC23 *KC31’’

g inc_mass3_2 sem tell

say ’’[cont]Next, I will explain it to you step by step. The rock’s
 kinetic energy at point A increases when the rock’s mass increases
 from 0.6 kg to 1 kg. The new value is $0.5 * 1 \text{ kg} * (v0.mag)^2 = 2 \text{ J}$
 while the old value is $0.5 * 0.6 \text{ kg} * (v0.mag)^2 = \{KE0. value\}$.
 ’’ kc ’’*KC20 *KC20q *KC23 *KC31’’
 say ’’[cont]In other words, the system’s new TME1 is the same as
 the old TME1, which is {TME1.value} but the system’s new TME0 is
 2 J, which is bigger than the old TME0, which is {TME0.value}.’’
 kc ’’*KC24 *KC23’’
 say ’’[cont] According to {eqn_CME01}, we have $W_{net01} = TME1 - TME0$.

When the rock's mass increased from 0.6 kg to 1 kg, the TME0 increased but TME1 remained the same, therefore, Wnet01 must decrease. That is, the net work done on the rock as it moves from A to B *decreases* as the rock's mass increases from 0.6 kg to 1 kg.''

```
g compare_path
do compare_path_1
do compare_path_2 rand sem justify
```

```
g compare_path_1 sem elicit
say ''[menu 'Yes, it would affect|No, it would not affect|
I do not know|not enough information|I do not know.']]Let's go
back to the original situation. The problem does not give any
hint about the path taken by the rock. Would taking different
paths affect our final answer of the net work done on the rock
from A to B? ''
if ''No, it would not affect.''' true kc ''+KC28 +KC24 +KC20''
otherwise say ''[cont]Actually ... taking different paths would not
affect our final answer of the net work done on the rock
from A to B. '' kc ''-KC28 -KC24 -KC20 *KC28 *KC24 *KC20''
```

```
g compare_path_1 sem tell
say ''[cont]Let's go back to the original situation. The problem does
not give any hint about the path taken by the rock. However,
taking different paths would not affect our final answer of the
net work done on the rock from A to B. ''
kc ''*KC28 *KC24 *KC20''
```

```

g compare_path_2 sem elicit
say ''Why? ''
if ''As long as the rock's mass, the rock's magnitude of v0 and v1 are
still the same, then Wnet01 is still {Wnet01.value}. ''
true kc ''+KC28 +KC24 +KC20 +KC15''
otherwise say ''[cont]This is because the path that the rock took
from point A to point B was not involved in the problem solving,
only the rock's mass, the rock's magnitude of velocity at T0 and
T1 were involved. Thus, as long as they are still the same,
Wnet01 is still {Wnet01.value}. ''
kc ''-KC28 -KC24 -KC20 *KC28 *KC24 *KC 20 *KC15''

```

```

g compare_path_2 sem tell
say ''[cont]This is because the path that the rock took from point
A to point B was not involved in the problem solving, only the
rock's mass, the rock's magnitude of velocity at T0 and T1 were
involved. Thus, as long as they are still the same, Wnet01 is
still {Wnet01.value}. '' kc ''*KC28 *KC24 *KC20 *KC15''

```

APPENDIX J

STUDY 2: TUTORIAL FEATURES

Table J1: Representing Sample Dialogue in Table 5.2 Using the 18 Feature Variables From Study 2

	29	30	31	32	33	34	35	36	37
Order ^a	S	T	S	T	S	T	S	T	S
Speaker ^b	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
TMOVE ^c	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
Features:									
v5durationBetweenDecisionT	0.03	nil	0	nil	0	nil	0	nil	0
TimeInSessionT	0.04	nil	0.04	nil	0.04	nil	0.04	nil	0.04
TimeBetweenSessionT	0	nil	0	nil	0	nil	0	nil	0
Rewards	0	0	0	0	0	0	0	0	0
v12EarlyTrainingPS	0	nil	0	nil	0	nil	0	nil	0
v13SimpleProblemPS	0.5	nil	0.5	nil	0.5	nil	0.5	nil	0.5
v14DuringWalkThroughPS	1	nil	1	nil	1	nil	1	nil	1
v15nKCsPS	0.12	nil	0.12	nil	0.15	nil	0.15	nil	0.15
v16nKCsSessionPS	0.28	nil	0.28	nil	0.3	nil	0.3	nil	0.3
v27tellsSinceElicitA	0	nil	0.08	nil	0	nil	0	nil	0.08
v28pctElicitA	0.5	nil	0.49	nil	0.51	nil	0.52	nil	0.51
v31pctTellsKCSessionA	0.62	nil	0.62	nil	0.64	nil	0.63	nil	0.63
v32pctCorrectPM	0.8	nil	0.8	nil	0.8	nil	0.8	nil	0.8
v33pctOverallCorrectPM	0.57	nil	0.57	nil	0.56	nil	0.56	nil	0.56
v34pctCorrectSessionPM	0.83	nil	0.83	nil	0.84	nil	0.84	nil	0.84
v40nCorrectKCPM	0.26	nil	0.26	nil	0.28	nil	0.3	nil	0.3
v42pctOverallCorrectKCPM	0.75	nil	0.75	nil	0.77	nil	0.78	nil	0.78
v44pctCorrectKCPM	0.51	nil	0.51	nil	0.49	nil	0.51	nil	0.51
v45pctCorrectKCSessionPM	0.31	nil	0.31	nil	0.3	nil	0.31	nil	0.31
v46nIncorrectKCPM	0.11	nil	0.11	nil	0.11	nil	0.11	nil	0.11

^aThe number of the decision orders in the tutorial dialogue on KC_{20}

^b“S:” refers to the current state when the tutor make the decision while “T:” means that it is tutor decision turn

^cThe action the tutor decided to take

APPENDIX K

STUDY 2: DICHGAIN TUTORIAL TACTICS

Policy: **KC-general policy on Justify/Skip-Justify**

Features: [18]

ECR: 0.572656

ECR Lower Bound: 0.19178

Mapping: [.3060]

States: ['0 skip', '1 else']

Policy: **KC-general policy on Elicit/Tell**

Features: [5,9,17,30]

ECR: 0.974232

ECR Lower Bound: 0.262842

Mapping: [11.0000,0.5,.5063,.6459]

States: ['1:1:1:1 tell', '0:0:1:1 elicit', '0:1:1:0 tell', '0:0:0:0 elicit', '1:1:1:0 tell', '0:1:0:0 elicit', '1:0:1:0 tell', '1:1:0:0 tell', '1:1:0:1 tell', '0:1:0:1 tell', '1:0:0:0 elicit', '1:0:1:1 elicit', '1:0:0:1 else', '0:1:1:1 elicit', '0:0:0:1 elicit', '0:0:1:0 elicit']

Policy: **1^aet^b**

Features: [10,13,16,29]

ECR: 20.196

ECR Lower Bound: 5.18551

Mapping: [None,162727.8586,.4961,.7333]

States: ['COMP:0:0:0 elicit', 'MED:1:0:0 elicit', 'COMP:1:0:0 tell', 'MED:1:1:0 tell', 'COMP:0:0:1 tell', 'COMP:1:1:0 tell', 'COMP:1:0:1 tell', 'MED:0:0:0 elicit', 'MED:1:0:1 elicit', 'COMP:0:1:0 elicit', 'MED:0:0:1 elicit', 'COMP:0:1:1 elicit', 'MED:1:1:1 tell', 'COMP:1:1:1 elicit', 'MED:0::1 elicit', 'MED:0:1:1 elicit', 'MED:0:1:0 elicit', 'MED:0::0 elicit']

^aKC number. In this case, it is a KC-specific policy on KC1. Same below

^bType of tutorial decisions. “et” refers to Elicit/Tell decisions while “skip” refers to Justify/Skip-Justify decisions

Policy: **3et**

Features: [9,12,16,17]

ECR: -25.2687

ECR Lower Bound: -45.363

Mapping: [0.5,1575.5233,.4980,.5558]

States: ['0:0:1:1 elicit', '0:1:1:0 elicit', '0:0:0:0 elicit', '0:1:0:0 elicit', '0:0::1 tell', '0:1:0:1 tell', '0:1:1:1 elicit', '0:0:0:1 elicit', '0:0:1:0 tell']

Policy: **5et**

Features: [9,13,15,25]

ECR: -4.57221

ECR Lower Bound: -13.6603

Mapping: [0.5,80963.9211,1.0000,2.0000]

States: ['1:1:1:1 elicit', '0:0:1:1 elicit', '0:1:1:0 tell', '0:0:0:0 elicit', '1:1:1:0 elicit', '0:1:0:0 tell', '1:0:1:0 elicit', '1:1:0:0 elicit', '1:1:0:1 elicit', '0:1:0:1 elicit', '1:0:0:0 elicit', '1:0:1:1 elicit', '1:0:0:1 elicit', '0:1:1:1 elicit', '0:0:0:1 tell', '0:0:1:0 elicit']

Policy: **9et**

Features: [12]

ECR: 63.9232

ECR Lower Bound: 58.7516

Mapping: [1056.5808]

States: ['0 tell', '1 else']

Policy: **12skip**

Features: [5,9,15,17]

ECR: 0

ECR Lower Bound: 0

Mapping: [65.0000,0.5,1.0000,.4422]

States: ['1:1:1:1 noskip', '0:1:1:0 noskip', '1:1:1:0 noskip', '0:1:0:0 noskip', '1:1:0:0 noskip', '1:1:0:1 noskip', '0:1:0:1 noskip', '0:1:0:1 noskip', '0:1:1:1 noskip']

Policy: **12et**

Features: [5,9,15,17]

ECR: 0

ECR Lower Bound: 0

Mapping: [65.0000,0.5,1.0000,.4422]

States: ['1:1:1:1 elicit', '0:1:1:0 elicit', '1:1:1:0 elicit', '0:1:0:0 elicit', '1:1:0:0 elicit', '1:1:0:1 elicit', '0:1:0:1 elicit', '0:1:1:1 elicit']

Policy: **13et**

Features: [5,9,15,28]

ECR: 14.5194

ECR Lower Bound: 1.08511

Mapping: [166.0000,0.5,1.0000,5.0000]

States: ['1:1:1:1 elicit', '0:0:1:1 elicit', '0:1:1:0 elicit', '0:0:0:0 elicit', '1:1:1:0 elicit', '0:1:0:0 elicit', '1:0:1:0 tell', '1:1:0:0 elicit', '1:1:0:1 elicit', '0:1:0:1 elicit', '1:0:0:0 elicit', '1:0:1:1 tell', '1:0:0:1 tell', '0:1:1:1 elicit', '0:0:0:1 elicit', '0:0:1:0 elicit']

Policy: **14skip**

Features: [5,9,15,18]

ECR: 17.6695

ECR Lower Bound: 14.503

Mapping: [94.0000,0.5,1.0000,.3517]

States: ['1:1:1:1 else', '0:0:1:1 else', '0:1:1:0 else', '0:0:0:0 noskip', '1:1:1:0 else', '0:1:0:0 else', '1:0:1:0 noskip', '1:1:0:0 else', '1:1:0:1 else', '0:1:0:1 else', '1:0:0:0 noskip', '1:0:1:1 noskip', '1:0:0:1 else', '0:1:1:1 else', '0:0:1:0 noskip']

Policy: **14et**

Features: [9]

ECR: 54.151

ECR Lower Bound: 47.8956

Mapping: [0.5]

States: ['0 tell', '1 elicit']

Policy: **15et**

Features: [5,9,15,17]

ECR: 0

ECR Lower Bound: 0

Mapping: [344.0000,0.5,1.0000,.4565]

States: ['1:1:1:1 elicit', '0:0:1:1 elicit', '0:1:1:0 elicit', '0:0:0:0 tell', '1:1:1:0 elicit', '0:1:0:0 elicit', '1:0:1:0 elicit', '1:1:0:0 elicit', '1:1:0:1 elicit', '0:1:0:1 elicit', '1:0:0:0 elicit', '1:0:1:1 elicit', '1:0:0:1 elicit', '0:1:1:1 elicit', '0:0:0:1 elicit', '0:0:1:0 elicit']

Policy: **17et**

Features: [5,9,15,17]

ECR: 3.27521

ECR Lower Bound: 0.574497

Mapping: [88.5000,0.5,1.0000,.5610]

States: ['0:0:1:1 else', '0:0:0:0 elicit', '1:0:1:0 elicit', '1:0:0:0 else', '1:0:1:1 elicit', '1:0:0:1 else', '0:0:0:1 elicit', '0:0:1:0 else']

Policy: **18et**

Features: [5,9,15,17]

ECR: 0

ECR Lower Bound: 0

Mapping: [1.0000,0.5,1.0000,.4578]

States: ['1:1:1:1 elicit', '0:1:1:0 elicit', '1:1:1:0 elicit', '1:1:0:0 elicit', '1:1:0:1 elicit', '0:1:1:1 elicit']

Policy: **20skip**

Features: [5,14,15,26]

ECR: 4.28925

ECR Lower Bound: 0.376654

Mapping: [52.0000,0.5,1.0000,.5405]

States: ['1:1:1:1 skip', '0:0:1:1 else', '0:1:1:0 else', '0:0:0:0 noskip', '1:1:1:0 noskip', '0:1:0:0
else', '1:0:1:0 else', '1:1:0:0 else', '1:1:0:1 skip', '0:1:0:1 else', '1:0:0:0 skip', '1:0:1:1
else', '1:0:0:1 skip', '0:1:1:1 else', '0:0:0:1 else', '0:0:1:0 noskip']

Policy: **20et**

Features: [5,10,15,26]

ECR: 4.80774

ECR Lower Bound: 0.753813

Mapping: [52.0000,None,1.0000,.5405]

States: ['0:SIM:1:0 tell', '1:SIM:1:1 tell', '1:MED:0:1 tell', '0:COMP:0:0
elicit', '1:MED:0:0 elicit', '1:COMP:1:0 elicit', '0:MED:1:0 elicit', '1:MED:1:0
elicit', '1:COMP:0:0 else', '1:MED:1:1 tell', '0:SIM:0:0 tell', '1:COMP:1:1
elicit', '0:COMP:1:1 elicit', '1:SIM:1:0 elicit', '0:MED:0:0 elicit', '0:COMP:1:0
else', '1:SIM:0:0 tell', '0:SIM:1:1 tell', '0:MED:1:1 tell', '0:SIM:0:1 tell', '0:MED:0:1
tell', '0:COMP:0:1 tell', '1:SIM:0:1 tell', '1:COMP:0:1 else']

Policy: **21skip**

Features: [5,10,15,18]

ECR: 17.958

ECR Lower Bound: 12.1813

Mapping: [50.0000,None,0.0001,.2990]

States: ['1:MED:0:1 else', '0:COMP:0:0 else', '1:MED:0:0 else', '1:COMP:1:0
else', '0:MED:1:0 else', '1:MED:1:0 noskip', '1:COMP:0:0 noskip', '1:MED:1:1
else', '1:COMP:1:1 skip', '0:COMP:1:1 skip', '0:MED:0:0 else', '0:COMP:1:0
skip', '0:MED:1:1 else', '0:MED:0:1 else', '0:COMP:0:1 skip', '1:COMP:0:1 else']

Policy: **21et**

Features: [5,10,15,29]

ECR: 15.4776

ECR Lower Bound: 7.84628

Mapping: [50.0000,None,0.0001,.7179]

States: ['1:MED:0:1 tell', '0:COMP:0:0 else', '1:MED:0:0 tell', '1:COMP:1:0
elicit', '0:MED:1:0 elicit', '1:MED:1:0 tell', '1:COMP:0:0 else', '1:MED:1:1
tell', '1:COMP:1:1 else', '0:COMP:1:1 elicit', '0:MED:0:0 elicit', '0:COMP:1:0
elicit', '0:MED:1:1 elicit', '0:MED:0:1 tell', '0:COMP:0:1 elicit', '1:COMP:0:1 elicit']

Policy: **22skip**

Features: [5,9,24,30]

ECR: 27.4539

ECR Lower Bound: 9.55924

Mapping: [56.0000,0.5,75.0000,.6452]

States: ['1:1:1:1 else', '0:0:1:1 skip', '0:1:1:0 else', '0:0:0:0 noskip', '1:1:1:0 else', '0:1:0:0
else', '1:0:1:0 noskip', '1:1:0:0 else', '1:1:0:1 else', '0:1:0:1 else', '1:0:0:0 else', '1:0:1:1
else', '1:0:0:1 else', '0:1:1:1 else', '0:0:0:1 noskip', '0:0:1:0 noskip']

Policy: **22et**

Features: [5,14,15,17]

ECR: 9.39791

ECR Lower Bound: -5.36517

Mapping: [56.0000,0.5,1.0000,.5120]

States: ['1:1:1:1 tell','0:0:1:1 elicit','0:1:1:0 elicit','0:0:0:0 tell','1:1:1:0 tell','0:1:0:0 elicit','1:0:1:0 elicit','1:1:0:0 tell','1:1:0:1 elicit','0:1:0:1 elicit','1:0:0:0 tell','1:0:1:1 elicit','1:0:0:1 tell','0:1:1:1 elicit','0:0:0:1 tell','0:0:1:0 else']

Policy: **23skip**

Features: [5,14,16,28]

ECR: 47.2221

ECR Lower Bound: 30.2939

Mapping: [96.0000,0.5,.4979,2.0000]

States: ['1:1:1:1 else','0:0:1:1 noskip','0:1:1:0 noskip','0:0:0:0 skip','1:1:1:0 noskip','0:1:0:0 noskip','1:0:1:0 else','1:1:0:0 noskip','1:1:0:1 noskip','0:1:0:1 noskip','1:0:0:0 else','1:0:1:1 else','1:0:0:1 noskip','0:1:1:1 noskip','0:0:0:1 skip','0:0:1:0 skip']

Policy: **23et**

Features: [5]

ECR: 42.4539

ECR Lower Bound: 22.8901

Mapping: [96.0000]

States: ['0 else', '1 tell']

Policy: **24skip**

Features: [5,9,15,18]

ECR: 2.92158

ECR Lower Bound: -0.302659

Mapping: [61.0000,0.5,1.0000,.2981]

States: ['0:0:1:1 skip','0:0:0:0 noskip','1:0:1:0 noskip','1:0:0:0 noskip','1:0:1:1
else','1:0:0:1 else','0:0:0:1 noskip','0:0:1:0 noskip']

Policy: **24et**

Features: [13,14,15,29]

ECR: 7.23353

ECR Lower Bound: 2.71676

Mapping: [91143.4986,0.5,1.0000,.7353]

States: ['1:1:1:1 tell','0:0:1:1 tell','0:1:1:0 tell','0:0:0:0 tell','1:1:1:0 elicit','0:1:0:0
else','1:0:1:0 elicit','1:1:0:0 else','1:1:0:1 tell','0:1:0:1 tell','1:0:0:0 else','1:0:1:1
elicit','1:0:0:1 tell','0:1:1:1 else','0:0:0:1 tell','0:0:1:0 else']

Policy: **25et**

Features: [5,10,29,30]

ECR: 22.5096

ECR Lower Bound: 7.75064

Mapping: [400.5000,None,.7273,.6347]

States: ['1:MED:0:1 tell','0:COMP:0:0 else','1:MED:0:0 tell','1:COMP:1:0
tell','0:MED:1:0 else','1:MED:1:0 tell','1:COMP:0:0 elicit','1:MED:1:1
tell','1:COMP:1:1 elicit','0:COMP:1:1 tell','0:MED:0:0 elicit','0:COMP:1:0
elicit','0:MED:1:1 else','0:MED:0:1 elicit','0:COMP:0:1 else','1:COMP:0:1 tell']

Policy: **26skip**

Features: [9,13,15,26]

ECR: 19.2677

ECR Lower Bound: 8.82795

Mapping: [0.5,166692.6890,1.0000,.3333]

States: ['0:0:1:1 noskip', '0:1:1:0 else', '0:0:0:0 else', '0:1:0:0 else', '0:1:0:1 noskip', '0:1:1:1 skip', '0:0:0:1 noskip', '0:0:1:0 else']

Policy: **26et**

Features: [5,9,15,26]

ECR: 18.7551

ECR Lower Bound: 10.592

Mapping: [70.0000,0.5,1.0000,.3333]

States: ['0:0:1:1 tell', '0:0:0:0 tell', '1:0:1:0 tell', '1:0:0:0 elicit', '1:0:1:1 elicit', '1:0:0:1 else', '0:0:0:1 else', '0:0:1:0 elicit']

Policy: **27skip**

Features: [9,13,15,17]

ECR: 10.5732

ECR Lower Bound: 2.7791

Mapping: [0.5,164567.3945,1.0000,.5406]

States: ['0:0:1:1 else', '0:1:1:0 noskip', '0:0:0:0 skip', '0:1:0:0 noskip', '0:1:0:1 else', '0:1:1:1 else', '0:0:0:1 else', '0:0:1:0 noskip']

Policy: **27et**

Features: [5,10,27,30]

ECR: 16.7804

ECR Lower Bound: 5.94975

Mapping: [74.0000,None,.3077,.6270]

States: ['1:MED:0:1 else','0:COMP:0:0 elicit','1:MED:0:0 elicit','1:COMP:1:0
tell','0:MED:1:0 tell','1:MED:1:0 tell','1:COMP:0:0 else','1:MED:1:1
tell','1:COMP:1:1 tell','0:COMP:1:1 elicit','0:MED:0:0 elicit','0:COMP:1:0
tell','0:MED:1:1 else','0:MED:0:1 elicit','0:COMP:0:1 elicit','1:COMP:0:1 else']

Policy: **28skip**

Features: [5,9,15,27]

ECR: 13.9743

ECR Lower Bound: 6.00922

Mapping: [59.0000,0.5,1.0000,.2083]

States: ['0:0:1:1 noskip','0:0:0:0 else','1:0:1:0 else','1:0:0:0 skip','1:0:1:1 else','1:0:0:1
skip','0:0:0:1 skip','0:0:1:0 else']

Policy: **28et**

Features: [9,13,16,27]

ECR: 15.2862

ECR Lower Bound: 2.52373

Mapping: [0.5,81595.8369,.4916,.2083]

States: ['0:0:1:1 else','0:1:1:0 tell','0:0:0:0 tell','0:1:0:0 tell','0:1:0:1 elicit','0:1:1:1
tell','0:0:0:1 elicit','0:0:1:0 elicit']

Policy: **31et**

Features: [5,9,15,17]

ECR: 1.29047

ECR Lower Bound: -8.00777

Mapping: [425.0000,0.5,0.0001,.5046]

States: ['1:1:1:1 elicit','0:0:1:1 elicit','0:1:1:0 tell','0:0:0:0 elicit','1:1:1:0 elicit','0:1:0:0 tell','1:0:1:0 else','1:1:0:0 tell','1:1:0:1 tell','0:1:0:1 elicit','1:0:0:0 elicit','1:0:1:1 tell','1:0:0:1 else','0:1:1:1 tell','0:0:0:1 elicit','0:0:1:0 else']

Policy: **32et**

Features: [9,12,26,30]

ECR: 14.5397

ECR Lower Bound: 0.175324

Mapping: [0.5,2387.6022,.4000,.6000]

States: ['0:0:1:1 tell','0:1:1:0 tell','0:0:0:0 elicit','0:1:0:0 tell','0:1:0:1 tell','0:1:1:1 tell','0:0:0:1 tell','0:0:1:0 elicit']

APPENDIX L

STUDY 3: EXAMPLE LOG 50 FEATURES

Representating the Sample Dialogue in Table 5.2 By 50 Feature Variables

Table L1: Student Autonomy Features

Order ^a :	29	30	31	32	33	34	35	36	37
Speaker ^b	S	T	S	T	S	T	S	T	S
TMOVE ^c	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v27tellsSinceElicitA	0	nil	0.08	nil	0	nil	0	nil	0.08
v28pctElicitA	0.5	nil	0.49	nil	0.51	nil	0.52	nil	0.51
v29stuWordsToTuWordsA	0.32	nil	0.32	nil	0.3	nil	0.31	nil	0.31
v30stuWordsToTuWordsSessionA	0.21	nil	0.21	nil	0.2	nil	0.2	nil	0.2
v31pctTellsKCSessionA	0.62	nil	0.62	nil	0.64	nil	0.63	nil	0.63

^aThe number of the decision orders in the tutorial dialogue on KC_{20}

^b“S:” refers to the current state when the tutor make the decision while “T:” means that it is tutor decision turn

^cThe action the tutor decided to take

Table L2: Problem Solving Context

Order:	29	30	31	32	33	34	35	36	37
Speaker	S	T	S	T	S	T	S	T	S
TMOVE	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v12EarlyTrainingPS	0	nil	0	nil	0	nil	0	nil	0
v13SimpleProblemPS	0.5	nil	0.5	nil	0.5	nil	0.5	nil	0.5
v14DuringWalkThroughPS	1	nil	1	nil	1	nil	1	nil	1
v15nKCsPS	0.12	nil	0.12	nil	0.15	nil	0.15	nil	0.15
v16nKCsSessionPS	0.28	nil	0.28	nil	0.3	nil	0.3	nil	0.3
v17newLevelDifficultyPS	1	nil	1	nil	1	nil	1	nil	1
v18conceptDifficultyPS	0.56	nil	0.5	nil	0.15	nil	0.15	nil	0.55
v19QuantativeDegreePS	0.5	nil	1	nil	0	nil	0	nil	0
v20numPhysConceptsTutorDialogueSessionPS	0.35	nil	0.35	nil	0.39	nil	0.39	nil	0.39
v21tutAverageConceptsPS	0.82	nil	0.82	nil	0.84	nil	0.84	nil	0.84
v22tutAverageConceptsSessionPS	0.66	nil	0.66	nil	0.68	nil	0.68	nil	0.68
v23tutConceptsToWordsPS	0.87	nil	0.87	nil	0.88	nil	0.88	nil	0.88
v24tutConceptsToWordsSessionPS	0.74	nil	0.74	nil	0.75	nil	0.75	nil	0.75
v25tutAverageWordsPS	0.85	nil	0.85	nil	0.86	nil	0.86	nil	0.86
v26tutAverageWordsSessionPS	0.75	nil	0.75	nil	0.76	nil	0.76	nil	0.76

Table L3: Background Features

Order:	29	30	31	32	33	34	35	36	37
Speaker	S	T	S	T	S	T	S	T	S
TMOVE	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v48genderBG	0	nil	0	nil	0	nil	0	nil	0
v49ageBG	0.38	nil	0.38	nil	0.38	nil	0.38	nil	0.38
v50MathSatBG	0.8	nil	0.8	nil	0.8	nil	0.8	nil	0.8
v51VerbalSatBG	0.67	nil	0.67	nil	0.67	nil	0.67	nil	0.67
v52pretestBG	0.69	nil	0.69	nil	0.69	nil	0.69	nil	0.69
v53averagePhysConceptsStudentDialogueSD	0.19	nil	0.19	nil	0.18	nil	0.19	nil	0.19

Table L4: Student Dialogue Features

Order:	29	30	31	32	33	34	35	36	37
Speaker	S	T	S	T	S	T	S	T	S
TMOVE	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v53averagePhysConceptsStudentDialogueSD	0.19	nil	0.19	nil	0.18	nil	0.19	nil	0.19
v54numStudentConceptualDialogueSD	0.23	nil	0.23	nil	0.23	nil	0.25	nil	0.25
v55stuConceptToWordRatioSD	0.1	nil	0.1	nil	0.1	nil	0.1	nil	0.1
v56stuAverageWordsSD	0.41	nil	0.41	nil	0.41	nil	0.42	nil	0.42
v57stuAverageConceptSD	0.31	nil	0.31	nil	0.3	nil	0.31	nil	0.31
v58averagePhysConceptsStudentDialogueSessionSD	0.19	nil	0.19	nil	0.18	nil	0.19	nil	0.19
v59numStudentConceptualDialogueSessionSD	0.44	nil	0.44	nil	0.44	nil	0.47	nil	0.47
v60stuConceptToWordRatioSessionSD	0.1	nil	0.1	nil	0.1	nil	0.1	nil	0.1
v61stuAverageWordsSessionSD	0.14	nil	0.14	nil	0.14	nil	0.14	nil	0.14
v62stuAverageConceptSessionSD	0.31	nil	0.31	nil	0.3	nil	0.31	nil	0.31

Table L5: Temporal Features

Order:	29	30	31	32	33	34	35	36	37
Speaker	S	T	S	T	S	T	S	T	S
TMOVE	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v5durationBetweenDecisionT	0.03	nil	0	nil	0	nil	0	nil	0
v9TimeInSessionT	0.04	nil	0.04	nil	0.04	nil	0.04	nil	0.04
v10TimeBetweenSessionT	0	nil	0	nil	0	nil	0	nil	0
Start	1187039153	nil	1187042098	nil	1187042117	nil	1187042212	nil	1187042252
End	1187042098	nil	1187042117	nil	1187042212	nil	1187042252	nil	1187042253

Table L6: Performance Features

Order:	29	30	31	32	33	34	35	36	37
Speaker	S	T	S	T	S	T	S	T	S
TMOVE	nil	tell	nil	elicit	nil	elicit_skippable	nil	tell	nil
v32pctCorrectPM	0.8	nil	0.8	nil	0.8	nil	0.8	nil	0.8
v33pctOverallCorrectPM	0.57	nil	0.57	nil	0.56	nil	0.56	nil	0.56
v34pctCorrectSessionPM	0.83	nil	0.83	nil	0.84	nil	0.84	nil	0.84
v35pctOverallCorrectSessionPM	0.31	nil	0.31	nil	0.3	nil	0.31	nil	0.31
KC	20	20	20	20	20	20	20	20	20
sem	nil	nil	nil	nil	nil	nil	nil	nil	nil
goal_name	nil	58-ke0_principle_selection	nil	58-ke0_write_eqn	nil	58-ke0_discuss_truth	nil	58-ke0_discuss_qualitative_1	nil
Justify	nil	nil	nil	nil	nil	nil	nil	nil	nil
v40nCorrectKCPM	0.26	nil	0.26	nil	0.28	nil	0.3	nil	0.3
v41nCorrectKCSessionPM	0.3	nil	0.3	nil	0.32	nil	0.32	nil	0.32
v42pctOverallCorrectKCPM	0.75	nil	0.75	nil	0.77	nil	0.78	nil	0.78
v43pctOverallCorrectKCSessionPM	0.83	nil	0.83	nil	0.84	nil	0.84	nil	0.84
v44pctCorrectKCPM	0.51	nil	0.51	nil	0.49	nil	0.51	nil	0.51
v45pctCorrectKCSessionPM	0.31	nil	0.31	nil	0.3	nil	0.31	nil	0.31
v46nIncorrectKCPM	0.11	nil	0.11	nil	0.11	nil	0.11	nil	0.11

APPENDIX M

STUDY 3: NORMGAIN AND INVNORMGAIN TUTORIAL TACTICS (FEATURES)

Table M1: NormGain Tutorial Tactics:

KC	action	#features	Corpus	Selection Method	ECR	Lower Bound	Upper Bound
KC1	ET	1	DichGain	0	36.33	13.76	44.11
Cols:	v24tutConceptsToWordsSessionPS						
KC14	ET	1	Comb	0	9.97	9.85	10.06
Cols:	v5durationBetweenDecisionT						
KC14	JS	1	Exp	0	10.89	10.89	10.89
Cols:	v12EarlyTrainingPS						
KC20	ET	3	Exp	ECR	14.25	10.04	18.12
Cols:	v18conceptDifficultyPS v23tutConceptsToWordsPS v26tutAverageWordsSessionPS						
KC20	JS	5	Exp	PCA-Hedge	20.23	15.11	22.12
Cols:	v9TimeInSessionT v15nKCsPS v28pctElicitA v56stuAverageWordsSD v62stuAverageConceptSessionSD						
KC21	JS	3	Exp	PCA-Hedge	32.26	29.8	34.59
Cols:	v13SimpleProblemPS v18conceptDifficultyPS v20numPhysConceptsTutorDialogueSessionPS						
KC21	ET	3	Exp	Upper_Bound	20.99	15.2	26.64
Cols:	v15nKCsPS v18conceptDifficultyPS v30stuWordsToTuWordsSessionA						
KC22	JS	5	Exp	Upper_Bound	37.02	22.45	47.12
Cols:	v23tutConceptsToWordsPS v24tutConceptsToWordsSessionPS v27tellsSinceElicitA v46nIncorrectKCPM v47nIncorrectKCSessionPM						
KC22	ET	2	Exp	Hedge	33.42	10.55	40.08
Cols:	v23tutConceptsToWordsPS v27tellsSinceElicitA						
KC24	JS	6	Exp	Upper_Bound	20.05	16.09	24.45
Cols:	v14DuringWalkThroughPS v18conceptDifficultyPS v49ageBG v59numStudentConceptualDialogueSessionSD v60stuConceptToWordRatioSessionSD v62stuAverageConceptSessionSD						
KC24	ET	4	DichGain	ECR	13.51	8.41	16.35
Cols:	v10TimeBetweenSessionT v14DuringWalkThroughPS v18conceptDifficultyPS v55stuConceptToWordRatioSD						
KC27	ET	4	Exp	PCA-Random	48.51	38.9	57
Cols:	v5durationBetweenDecisionT v17newLevelDifficultyPS v23tutConceptsToWordsPS v35pctOverallCorrectSessionPM						
KC27	JS	4	Exp	ECR	55.68	30.89	44.05
Cols:	v9TimeInSessionT v25tutAverageWordsPS v27tellsSinceElicitA v56stuAverageWordsSD						
KC28	JS	3	DichGain	Upper_Bound	30.11	15.82	41
Cols:	v15nKCsPS v24tutConceptsToWordsSessionPS v54numStudentConceptualDialogueSD						
KC28	ET	5	DichGain	Upper_Bound	20.45	8.45	30.43
Cols:	v5durationBetweenDecisionT v16nKCsSessionPS v18conceptDifficultyPS v29stuWordsToTuWordsA v41nCorrectKCSessionPM						
KC-general	ET	4	Exp	Lower_Bound	4.51	1.38	8.33
Cols:	v5durationBetweenDecisionT v18conceptDifficultyPS v27tellsSinceElicitA v46nIncorrectKCPM						
KC-general	JS	5	DichGain	ECR	9.84	6.77	12.19
Cols:	v16nKCsSessionPS v17newLevelDifficultyPS v23tutConceptsToWordsPS v25tutAverageWordsPS v27tellsSinceElicitA						

Table M2: InvNormGain Tutorial Tactics:

KC	action	#features	Corpus	Selection Method	ECR	Lower Bound	Upper Bound
KC0	JS	4	DichGain	ECR	28.7	21.91	33.91
Columns:	v17newLevelDifficultyPS		v23tutConceptsToWordsPS		v25tutAverageWordsPS		
	v27tellsSinceElicitA						
KC1	ET	2	DichGain	Upper_Bound	70.39	62.26	74.12
Cols:	v19QuantativeDegreePS v41nCorrectKCSessionPM						
KC14	ET	1	DichGain	0	87.63	87.63	87.63
Cols:	v10TimeBetweenSessionT						
KC14	JS	1	Comb	0	86.67	86.67	86.67
Cols:	v12EarlyTrainingPS						
KC20	ET	3	Exp	PCA-Lower_Bound	31.15	26.91	34.98
Cols:	v19QuantativeDegreePS v23tutConceptsToWordsPS v25tutAverageWordsPS						
KC20	JS	5	DichGain	Random	41.43	33.99	47.57
Cols:	v13SimpleProblemPS		v14DuringWalkThroughPS		v18conceptDifficultyPS		v44pctCorrectKCPM
	v55stuConceptToWordRatioSD						
KC21	ET	6	Exp	Hedge	23.43	18.01	27.06
Cols:	v5durationBetweenDecisionT		v14DuringWalkThroughPS		v19QuantativeDegreePS		
	v55stuConceptToWordRatioSD v56stuAverageWordsSD v61stuAverageWordsSessionSD						
KC21	JS	3	Exp	ECR	46.74	41.83	51.55
Cols:	v18conceptDifficultyPS v35pctOverallCorrectSessionPM v45pctCorrectKCSessionPM						
KC22	JS	4	DichGain	ECR	95.36	88.01	103.16
Cols:	v5durationBetweenDecisionT v19QuantativeDegreePS v27tellsSinceElicitA v32pctCorrectPM						
KC22	ET	3	Exp	ECR	79.29	58.21	102.28
Cols:	v24tutConceptsToWordsSessionPS v26tutAverageWordsSessionPS v27tellsSinceElicitA						
KC24	JS	2	DichGain	Upper_Bound	49.19	33.14	61.55
Cols:	v17newLevelDifficultyPS v47nIncorrectKCSessionPM						
KC24	ET	2	DichGain	ECR	36.36	30.77	42.19
Cols:	v14DuringWalkThroughPS v18conceptDifficultyPS						
KC27	ET	4	DichGain	Lower_Bound	37.7	28.76	45.02
Cols:	v5durationBetweenDecisionT v9TimeInSessionT v19QuantativeDegreePS v27tellsSinceElicitA						
KC27	JS	5	DichGain	Hedge	60.04	40.64	50.87
Cols:	v9TimeInSessionT		v40nCorrectKCPM		v42pctOverallCorrectKCPM		v46nIncorrectKCPM
	v56stuAverageWordsSD						
KC28	ET	3	DichGain	ECR	67.97	51.86	81.84
Cols:	v5durationBetweenDecisionT v18conceptDifficultyPS v46nIncorrectKCPM						
KC28	JS	5	Exp	Upper_Bound	67.37	35.71	58.29
Cols:	v14DuringWalkThroughPS		v15nKCsPS		v25tutAverageWordsPS		v46nIncorrectKCPM
	v60stuConceptToWordRatioSessionSD						
KC-general	ET	5	DichGain	Upper_Bound	9.28	5.94	12.54
Cols:	v16nKCsSessionPS		v18conceptDifficultyPS		v23tutConceptsToWordsPS		v31pctTellsKCSessionA
	v35pctOverallCorrectSessionPM						
KC-general	ET	5	DichGain	Upper_Bound	9.28	5.94	12.54
Cols:	v16nKCsSessionPS		v18conceptDifficultyPS		v23tutConceptsToWordsPS		v31pctTellsKCSessionA
	v35pctOverallCorrectSessionPM						

APPENDIX N

STUDY 3: NORMAGAIN TUTORIAL TACTICS (POLICIES)

Policy: 'KC-general policy on Elicit/Tell'

Features: [5,18,27,46]

ECR: 4.505940

ECR Lower Bound: 1.381010

Mapping: [[22.00581],[0.29775,0.60095],[2.3034],[88.0896,143.231]]

States: ['0:0:0:0 elicit','0:0:0:1 tell','0:0:0:2 elicit','0:0:1:0 elicit','0:0:1:1 tell','0:0:1:2 tell','0:1:0:0 elicit','0:1:0:1 tell','0:1:0:2 tell','0:1:1:0 elicit','0:1:1:1 else','0:1:1:2 tell','0:2:0:0 elicit','0:2:0:1 tell','0:2:0:2 else','0:2:1:0 elicit','0:2:1:1 else','0:2:1:2 else','1:0:0:0 tell','1:0:0:1 tell','1:0:0:2 tell','1:0:1:0 elicit','1:0:1:1 tell','1:0:1:2 tell','1:1:0:0 elicit','1:1:0:1 tell','1:1:0:2 tell','1:1:1:0 elicit','1:1:1:1 else','1:1:1:2 tell','1:2:0:0 elicit','1:2:0:1 tell','1:2:0:2 elicit','1:2:1:0 elicit','1:2:1:1 tell','1:2:1:2 tell']

Policy: 'KC-general policy on Justify/Skip-Justify'

Features: [16,17,23,25,27]

ECR: 9.844580

ECR Lower Bound: 6.774750

Mapping: [[92.26325,183.4505],[0.5],[0.075178],[22.759655,23.592848],[1.0]]

States: ['0:0:0:0 noskip', '0:0:0:1 noskip', '0:0:1:0 skip', '0:0:1:1 noskip', '0:0:2:0 noskip', '0:0:2:1 skip', '0:0:1:0:0 noskip', '0:0:1:0:1 noskip', '0:0:1:1:0 noskip', '0:0:1:1:1 noskip', '0:0:1:2:0 noskip', '0:0:1:2:1 skip', '0:1:0:0:0 skip', '0:1:0:0:1 noskip', '0:1:0:1:0 noskip', '0:1:0:1:1 skip', '0:1:0:2:0 noskip', '0:1:0:2:1 noskip', '0:1:1:0:0 skip', '0:1:1:0:1 noskip', '0:1:1:1:0 noskip', '0:1:1:1:1 noskip', '0:1:1:2:0 noskip', '0:1:1:2:1 skip', '1:0:0:0:0 noskip', '1:0:0:0:1 noskip', '1:0:0:1:0 noskip', '1:0:0:1:1 noskip', '1:0:0:2:0 skip', '1:0:0:2:1 noskip', '1:0:1:0:0 noskip', '1:0:1:0:1 noskip', '1:0:1:1:0 noskip', '1:0:1:1:1 noskip', '1:0:1:2:0 noskip', '1:0:1:2:1 noskip', '1:1:0:0:0 skip', '1:1:0:0:1 skip', '1:1:0:1:0 noskip', '1:1:0:1:1 noskip', '1:1:0:2:0 skip', '1:1:0:2:1 noskip', '1:1:1:0:0 noskip', '1:1:1:0:1 noskip', '1:1:1:1:0 noskip', '1:1:1:1:1 skip', '1:1:1:2:0 noskip', '1:1:1:2:1 noskip', '2:0:0:0:0 noskip', '2:0:0:0:1 skip', '2:0:0:1:0 noskip', '2:0:0:1:1 skip', '2:0:0:2:0 noskip', '2:0:0:2:1 noskip', '2:0:1:1:0 noskip', '2:0:1:1:1 skip', '2:0:1:2:0 noskip', '2:0:1:2:1 skip', '2:1:0:0:0 skip', '2:1:0:1:0 skip', '2:1:0:1:1 skip', '2:1:0:2:0 noskip', '2:1:0:2:1 skip', '2:1:1:1:0 skip', '2:1:1:1:1 noskip', '2:1:1:2:0 noskip', '2:1:1:2:1 skip']

Policy: 1^aet^b

Features: [24]

ECR: 36.331200

ECR Lower Bound: 13.755700

Mapping: [[0.062402,0.071909]]

States: ['0 elicit', '1 tell', '2 elicit']

^aKC number. In this case, means it is a KC-specific policy on KC1. Same below

^bType of tutorial decisions. "et" refers to Elicit/Tell decisions while "skip" refers to Justify/Skip-Justify decisions

Policy: **14et'**

Features: [5]

ECR: 9.966550

ECR Lower Bound: 9.848810

Mapping: [[160.073973]]

States: ['0 elicit', '1 else']

Policy: **14skip'**

Features: [12]

ECR: 10.887100

ECR Lower Bound: 10.887100

Mapping: [[0.5]]

States: ['0 skip', '1 noskip']

Policy: **20et'**

Features: [18,23,26]

ECR: 14.254900

ECR Lower Bound: 10.036400

Mapping: [[0.3778],[0.073981],[22.577706]]

States: ['0:0:0 elicit', '0:0:1 elicit', '0:1:0 tell', '0:1:1 else', '1:0:0 else', '1:0:1 elicit', '1:1:0 elicit', '1:1:1 elicit']

Policy: **20skip'**

Features: [9,15,28,56,62]

ECR: 20.227600

ECR Lower Bound: 15.113600

Mapping: [[3040.799131],[65.0135],[0.493592],[4.176873],[0.28604]]

States: ['0:0:0:0 noskip','0:0:0:1 skip','0:0:1:0 skip','0:0:1:1 skip','0:1:0:0 noskip','0:1:0:1 skip','0:1:1:0 noskip','0:1:1:1 skip','0:1:0:0 skip','0:1:0:1 noskip','0:1:1:0 skip','0:1:1:1 noskip','1:0:0:0 noskip','1:0:0:1 skip','1:0:1:0 noskip','1:0:1:1 skip','1:1:0:0 skip','1:1:0:1 noskip','1:1:1:0 skip','1:1:1:1 noskip']

Policy: **21et'**

Features: [15,18,30]

ECR: 20.993900

ECR Lower Bound: 15.201400

Mapping: [[44.15375],[0.30125,0.6075],[0.08421]]

States: ['0:0:0 elicit','0:0:1 elicit','0:1:0 elicit','0:1:1 elicit','0:2:0 elicit','0:2:1 elicit','1:0:0 elicit','1:0:1 elicit','1:1:0 elicit','1:1:1 tell','1:2:0 else','1:2:1 else']

Policy: **21skip'**

Features: [13,18,20]

ECR: 32.255300

ECR Lower Bound: 29.801900

Mapping: [[0.75],[0.30125,0.6075],[235.1497,468.0707]]

States: ['0:0:0 noskip','0:0:1 noskip','0:0:2 skip','0:1:0 else','0:1:1 else','0:1:2 else','0:2:0 noskip','0:2:1 noskip','0:2:2 noskip','1:0:0 else','1:0:1 skip','1:0:2 else','1:1:0 noskip','1:1:1 noskip','1:1:2 noskip','1:2:0 skip','1:2:1 skip','1:2:2 skip']

Policy: **22et'**

Features: [23,27]

ECR: 33.419600

ECR Lower Bound: 10.554100

Mapping: [[0.067549,0.07119,0.074431],[2.2836]]

States: ['0:0 elicit','0:1 tell','1:0 tell','1:1 else','2:0 tell','2:1 else','3:0 tell','3:1 else']

Policy: **22skip'**

Features: [23,24,27,46,47]

ECR: 37.015500

ECR Lower Bound: 22.445200

Mapping: [[0.067549,0.07119,0.074431],[0.067915],[2.2836],[4.197,7.44825],[9.085]]

States: ['0:0:0:0 noskip','0:0:0:1 else','0:0:1:0 else','0:0:1:1 else','0:0:1:0:0
else','0:0:1:0:1 else','0:0:1:1:0 else','0:0:1:1:1 else','0:0:1:2:0 else','0:1:0:0:0
else','0:1:1:0:0 else','1:0:0:0:0 noskip','1:0:0:0:1 else','1:0:0:1:0 skip','1:0:0:1:1
else','1:0:0:2:0 skip','1:0:0:2:1 else','1:0:1:0:0 noskip','1:0:1:0:1 else','1:0:1:1:0
else','1:0:1:1:1 skip','1:0:1:2:0 noskip','1:0:1:2:1 else','1:1:0:0:0 else','1:1:0:0:1
else','1:1:0:1:0 skip','1:1:0:1:1 noskip','1:1:0:2:0 else','1:1:0:2:1 skip','1:1:1:0:0
else','1:1:1:0:1 else','1:1:1:1:0 else','1:1:1:1:1 noskip','1:1:1:2:1 else','2:0:0:0:0
skip','2:0:0:0:1 else','2:0:0:1:0 else','2:0:0:1:1 noskip','2:0:0:2:0 else','2:0:0:2:1
skip','2:0:1:0:0 else','2:0:1:1:0 noskip','2:0:1:1:1 else','2:0:1:2:0 noskip','2:0:1:2:1
noskip','2:1:0:0:0 skip','2:1:0:0:1 else','2:1:0:1:0 noskip','2:1:0:1:1 else','2:1:0:2:0
else','2:1:0:2:1 noskip','2:1:1:0:0 else','2:1:1:0:1 else','2:1:1:1:0 skip','2:1:1:1:1
noskip','2:1:1:2:0 else','2:1:1:2:1 else','3:0:0:0:0 else','3:0:0:1:0 else','3:0:0:2:0
else','3:0:0:2:1 else','3:0:1:0:0 else','3:0:1:2:0 else','3:0:1:2:1 else','3:1:0:0:0
else','3:1:0:0:1 skip','3:1:0:1:0 else','3:1:0:1:1 else','3:1:0:2:0 else','3:1:0:2:1
else','3:1:1:0:0 else','3:1:1:0:1 noskip','3:1:1:1:0 else','3:1:1:1:1 else','3:1:1:2:0
else','3:1:1:2:1 skip']

Policy: **24et'**

Features: [10,14,18,55]

ECR: 13.511000

ECR Lower Bound: 8.405400

Mapping: [[218512.016452],[0.5],[0.52575],[0.07082,0.086168,0.106655]]

States: ['0:0:0:0 else','0:0:0:1 elicit','0:0:0:2 elicit','0:0:0:3 elicit','0:0:1:0 else','0:0:1:1 tell','0:0:1:2 tell','0:0:1:3 tell','0:1:0:0 else','0:1:0:1 else','0:1:0:2 else','0:1:0:3 elicit','0:1:1:0 else','0:1:1:1 else','0:1:1:2 else','0:1:1:3 else','1:0:0:0 elicit','1:0:0:1 elicit','1:0:0:2 else','1:0:0:3 elicit','1:0:1:0 elicit','1:0:1:1 else','1:0:1:2 elicit','1:0:1:3 tell','1:1:0:0 elicit','1:1:0:1 tell','1:1:0:2 elicit','1:1:0:3 elicit','1:1:1:0 else','1:1:1:1 else','1:1:1:2 else','1:1:1:3 else']

Policy: **24skip'**

Features: [14,18,49,59,60,62]

ECR: 20.046000

ECR Lower Bound: 16.093400

Mapping: [[0.5],[0.341425],[20.999997],[14.0911],[0.1],[0.266733]]

States: ['0:0:0:0:0 noskip', '0:0:0:0:1:0 skip', '0:0:0:1:0:0 skip', '0:0:0:1:0:1 skip',
skip', '0:0:0:1:1:0 noskip', '0:0:0:1:1:1 skip', '0:0:1:0:0:0 noskip', '0:0:1:0:1:0 skip',
skip', '0:0:1:1:0:0 skip', '0:0:1:1:0:1 skip', '0:0:1:1:1:0 skip', '0:0:1:1:1:1 skip',
skip', '0:1:0:0:0:0 skip', '0:1:0:0:0:1 noskip', '0:1:0:0:1:0 skip', '0:1:0:0:1:1 skip',
skip', '0:1:0:1:0:0 noskip', '0:1:0:1:0:1 noskip', '0:1:0:1:1:0 noskip', '0:1:0:1:1:1 noskip',
skip', '0:1:1:0:0:0 noskip', '0:1:1:0:0:1 noskip', '0:1:1:0:1:0 noskip', '0:1:1:0:1:1 noskip',
skip', '0:1:1:1:0:0 skip', '0:1:1:1:0:1 noskip', '0:1:1:1:1:0 skip', '0:1:1:1:1:1 skip',
skip', '1:0:0:0:0:0 skip', '1:0:0:0:0:1 skip', '1:0:0:0:1:0 skip', '1:0:0:0:1:1 skip', '1:0:0:1:0:0 skip',
noskip', '1:0:0:1:0:1 noskip', '1:0:0:1:1:0 noskip', '1:0:0:1:1:1 skip', '1:0:1:0:0:0 skip',
skip', '1:0:1:0:0:1 skip', '1:0:1:0:1:0 skip', '1:0:1:0:1:1 skip', '1:0:1:1:0:0 noskip',
skip', '1:0:1:1:0:1 noskip', '1:0:1:1:1:0 skip', '1:0:1:1:1:1 skip', '1:1:0:0:0:0 skip',
skip', '1:1:0:0:0:1 skip', '1:1:0:0:1:0 noskip', '1:1:0:0:1:1 noskip', '1:1:0:1:0:0 noskip',
noskip', '1:1:0:1:0:1 noskip', '1:1:0:1:1:0 noskip', '1:1:0:1:1:1 noskip', '1:1:1:0:0:0 noskip',
skip', '1:1:1:0:0:1 skip', '1:1:1:0:1:0 noskip', '1:1:1:0:1:1 noskip', '1:1:1:0:1:1 noskip',
skip', '1:1:1:0:1:1 skip', '1:1:1:1:0:0 skip', '1:1:1:1:0:1 noskip', '1:1:1:1:1:0 skip',
noskip', '1:1:1:1:1:0 noskip', '1:1:1:1:1:1 skip', '1:1:1:1:1:1 noskip']

Policy: **27et'**

Features: [5,17,23,35]

ECR: 48.509500

ECR Lower Bound: 38.899400

Mapping: [[78.989305],[0.5],[0.072768],[0.348344]]

States: ['0:0:0:0 else', '0:0:0:1 else', '0:0:1:0 else', '0:0:1:1 else', '0:1:0:0 else', '0:1:0:1 else',
tell', '0:1:1:0 else', '0:1:1:1 tell', '1:0:0:0 tell', '1:0:0:1 tell', '1:0:1:0 elicit', '1:0:1:1 elicit',
tell', '1:1:0:0 tell', '1:1:0:1 else', '1:1:1:0 else', '1:1:1:1 else']

Policy: **27skip'**

Features: [9,25,27,56]

ECR: 55.684600

ECR Lower Bound: 30.887000

Mapping: [[3657.04962],[22.760564,23.335972],[0.999999],[4.203393]]

States: ['0:0:0:0 skip','0:0:0:1 skip','0:0:1:0 skip','0:0:1:1 noskip','0:1:0:0 noskip','0:1:0:1 noskip','0:1:1:0 noskip','0:1:1:1 noskip','0:2:0:0 skip','0:2:0:1 skip','0:2:1:0 skip','0:2:1:1 noskip','1:0:0:0 noskip','1:0:0:1 noskip','1:0:1:0 skip','1:0:1:1 skip','1:1:0:0 skip','1:1:0:1 noskip','1:1:1:0 skip','1:1:1:1 noskip','1:2:0:0 skip','1:2:0:1 noskip','1:2:1:0 noskip','1:2:1:1 noskip']

Policy: **28et'**

Features: [5,16,18,29,41]

ECR: 20.446500

ECR Lower Bound: 8.450380

Mapping: [[73.053098],[114.40575],[0.393393],[0.046351],[32.2232]]

States: ['0:0:0:0:0 tell','0:0:0:1:0 elicit','0:0:1:0:0 else','0:0:1:1:0 elicit','0:1:0:0:0 else','0:1:0:0:1 tell','0:1:0:1:0 else','0:1:0:1:1 else','0:1:1:0:0 else','0:1:1:0:1 tell','0:1:1:1:0 elicit','0:1:1:1:1 tell','1:0:0:0:0 elicit','1:0:0:1:0 tell','1:0:1:0:0 else','1:0:1:1:0 tell','1:0:1:1:0 elicit','1:0:1:1:1 tell','1:1:0:0:0 else','1:1:0:0:1 tell','1:1:0:1:0 elicit','1:1:0:1:1 tell','1:1:1:0:0 else','1:1:1:0:1 tell','1:1:1:1:0 tell','1:1:1:1:1 tell']

Policy: **28skip'**

Features: [15,24,54]

ECR: 30.114700

ECR Lower Bound: 15.824200

Mapping: [[20.42295],[0.075692],[21.0157]]

States: ['0:0:0 noskip','0:0:1 noskip','0:1:0 noskip','0:1:1 noskip','1:0:0 noskip','1:0:1 skip','1:1:0 noskip','1:1:1 skip']

APPENDIX O

STUDY 3: INVNORMGAIN TUTORIAL TACTICS (POLICIES)

Policy: 'KC-general policy on Elicit/Tell'

Features: [16,18,23,31,35]

ECR: 9.282570

ECR Lower Bound: 5.942560

Mapping: [[113.0632],[0.2696,0.60035],[0.075261],[0.2993,0.61905,0.78495],[0.2504]]

States: ['0:0:0:0 elicit','0:0:0:1 tell','0:0:0:1:0 elicit','0:0:0:1:1 elicit','0:0:0:2:0 elicit','0:0:0:2:1 elicit','0:0:0:3:0 elicit','0:0:1:0:0 elicit','0:0:1:0:1 tell','0:0:1:1:0 elicit','0:0:1:1:1 elicit','0:0:1:2:0 elicit','0:0:1:2:1 elicit','0:0:1:3:0 elicit','0:1:0:0:0 elicit','0:1:0:0:1 tell','0:1:0:1:0 else','0:1:0:1:1 tell','0:1:0:2:0 elicit','0:1:0:2:1 else','0:1:0:3:0 elicit','0:1:1:0:0 elicit','0:1:1:0:1 tell','0:1:1:1:0 tell','0:1:1:1:1 tell','0:1:1:2:0 elicit','0:1:1:2:1 elicit','0:1:1:3:0 elicit','0:2:0:0:0 tell','0:2:0:0:1 elicit','0:2:0:1:0 elicit','0:2:0:1:1 elicit','0:2:0:2:0 elicit','0:2:0:2:1 elicit','0:2:0:3:0 elicit','0:2:1:0:1 tell','0:2:1:1:0 elicit','0:2:1:1:1 elicit','0:2:1:2:0 elicit','0:2:1:2:1 elicit','0:2:1:3:0 else','1:0:0:1:0 elicit','1:0:0:1:1 elicit','1:0:0:2:0 elicit','1:0:0:2:1 elicit','1:0:0:3:0 elicit','1:0:1:1:0 elicit','1:0:1:1:1 tell','1:0:1:2:0 elicit','1:0:1:2:1 elicit','1:0:1:3:0 tell','1:1:0:1:0 elicit','1:1:0:1:1 tell','1:1:0:2:0 else','1:1:0:2:1 tell','1:1:0:3:0 elicit','1:1:1:1:0 elicit','1:1:1:1:1 tell','1:1:1:2:0 elicit','1:1:1:2:1 elicit','1:1:1:3:0 tell','1:2:0:1:0 elicit','1:2:0:1:1 tell','1:2:0:2:0 else','1:2:0:2:1 tell','1:2:0:3:0 elicit','1:2:1:1:0 elicit','1:2:1:1:1 tell','1:2:1:2:0 elicit','1:2:1:2:1 tell','1:2:1:3:0 elicit']

Policy: 'KC-general policy on Justify/Skip-Justify'

Features: [17,23,25,27]

ECR: 28.701200

ECR Lower Bound: 21.911800

Mapping: [[0.5],[0.075178],[22.759655,23.592848],[1.0]]

States: ['0:0:0:0 noskip', '0:0:0:1 noskip', '0:0:1:0 noskip', '0:0:1:1 noskip', '0:0:2:0 noskip', '0:0:2:1 noskip', '0:1:0:0 noskip', '0:1:0:1 noskip', '0:1:1:0 noskip', '0:1:1:1 noskip', '0:1:2:0 noskip', '0:1:2:1 skip', '1:0:0:0 skip', '1:0:0:1 noskip', '1:0:1:0 noskip', '1:0:1:1 skip', '1:0:2:0 noskip', '1:0:2:1 skip', '1:1:0:0 skip', '1:1:0:1 noskip', '1:1:1:0 noskip', '1:1:1:1 noskip', '1:1:2:0 noskip', '1:1:2:1 skip']

Policy: '1^aet^b'

Features: [19,41]

ECR: 70.394200

ECR Lower Bound: 62.264200

Mapping: [[0.066964],[22.0]]

States: ['0:0 elicit', '0:1 tell', '1:0 tell', '1:1 tell']

^aKC number. In this case, it is a KC-specific policy on KC1. Same below

^bType of tutorial decisions. "et" refers to Elicit/Tell decisions while "skip" refers to Justify/Skip-Justify decisions

Policy: '14et'

Features: [10]

ECR: 87.631600

ECR Lower Bound: 87.631600

Mapping: [[0.0]]

States: ['0 tell', '1 tell']

Policy: '14skip'

Features: [12]

ECR: 86.666700

ECR Lower Bound: 86.666700

Mapping: [[0.5]]

States: ['0 noskip', '1 noskip']

Policy: '20et'

Features: [19,23,25]

ECR: 31.148800

ECR Lower Bound: 26.913700

Mapping: [[0.22755,0.71735],[0.073981],[22.954695]]

States: ['0:0:0 elicit', '0:0:1 elicit', '0:1:0 elicit', '0:1:1 elicit', '1:0:0 else', '1:0:1 else', '1:1:0 elicit', '1:1:1 elicit', '2:0:0 tell', '2:0:1 tell', '2:1:0 elicit', '2:1:1 elicit']

Policy: '20skip'

Features: [13,14,18,44,55]

ECR: 41.425700

ECR Lower Bound: 33.987700

Mapping: [[0.75],[0.5],[0.530107],[0.133398],[0.07486,0.10008]]

States: ['0:0:0:0:0 noskip', '0:0:0:0:1 noskip', '0:0:0:0:2 noskip', '0:0:0:1:0 noskip', '0:0:0:1:1 noskip', '0:0:0:1:2 skip', '0:0:1:0:0 skip', '0:0:1:0:1 noskip', '0:0:1:0:2 noskip', '0:0:1:1:0 skip', '0:0:1:1:1 noskip', '0:0:1:1:2 skip', '0:1:0:0:0 skip', '0:1:0:0:1 noskip', '0:1:0:0:2 skip', '0:1:0:1:0 skip', '0:1:0:1:1 skip', '0:1:0:1:2 skip', '0:1:1:0:0 noskip', '0:1:1:0:1 noskip', '0:1:1:0:2 skip', '0:1:1:1:0 noskip', '0:1:1:1:1 noskip', '0:1:1:1:2 noskip', '1:0:1:0:0 noskip', '1:0:1:0:1 noskip', '1:0:1:0:2 noskip', '1:0:1:1:0 skip', '1:0:1:1:1 skip', '1:0:1:1:2 noskip', '1:1:0:0:0 noskip', '1:1:0:0:1 skip', '1:1:0:0:2 skip', '1:1:0:1:0 noskip', '1:1:0:1:1 skip', '1:1:0:1:2 noskip']

Policy: **21et**

Features: [5,14,19,55,56,61]

ECR: 23.429100

ECR Lower Bound: 18.009100

Mapping: [[50.00604],[0.5],[0.2361,0.7172],[0.079174,0.102734],[3.92965,5.668881],[5.8023]]

States: ['0:0:0:0:0 else', '0:0:0:0:1:0 elicit', '0:0:0:0:1:1 else', '0:0:0:0:2:1 elicit',
'0:0:0:1:0:0 elicit', '0:0:0:1:1:0 elicit', '0:0:0:1:1:1 else', '0:0:0:1:2:0 elicit',
'0:0:0:1:2:1 elicit', '0:0:0:2:0:0 elicit', '0:0:0:2:1:0 elicit', '0:0:0:2:1:1 elicit',
'0:0:0:2:2:0 elicit', '0:0:0:2:2:1 else', '0:0:1:0:0:0 else', '0:0:1:0:1:0 else',
'0:0:1:0:1:1 elicit', '0:0:1:0:2:1 else', '0:0:1:1:0:0 else', '0:0:1:1:1:0 else',
'0:0:1:1:1:1 elicit', '0:0:1:1:2:0 elicit', '0:0:1:1:2:1 elicit', '0:0:1:2:0:0 else',
'0:0:1:2:1:0 elicit', '0:0:1:2:2:0 else', '0:0:1:2:2:1 else', '0:0:2:0:0:0 elicit',
'0:0:2:0:1:0 elicit', '0:0:2:0:1:1 elicit', '0:0:2:0:2:1 elicit', '0:0:2:1:0:0 elicit',
'0:0:2:1:1:0 elicit', '0:0:2:1:1:1 tell', '0:0:2:1:2:0 elicit', '0:0:2:1:2:1 tell',
'0:0:2:2:0:0 elicit', '0:0:2:2:1:0 tell', '0:0:2:2:1:1 elicit', '0:0:2:2:2:0 elicit',
'0:1:0:0:0:0 else', '0:1:0:0:1:0 tell', '0:1:0:0:1:1 elicit', '0:1:0:0:2:0 elicit',
'0:1:0:0:2:1 elicit', '0:1:0:1:0:0 elicit', '0:1:0:1:0:1 tell', '0:1:0:1:1:0 elicit',
'0:1:0:1:1:1 tell', '0:1:0:1:2:0 else', '0:1:0:1:2:1 tell', '0:1:0:2:0:0 elicit',
'0:1:0:2:1:0 elicit', '0:1:0:2:1:1 elicit', '0:1:0:2:2:0 else', '0:1:0:2:2:1 tell',
'0:1:1:0:0:0 tell', '0:1:1:0:1:0 elicit', '0:1:1:0:1:1 else', '0:1:1:0:2:0 elicit',
'0:1:1:0:2:1 tell', '0:1:1:1:0:0 tell', '0:1:1:1:1:0 else', '0:1:1:1:1:1 elicit',
'0:1:1:1:2:0 tell', '0:1:1:1:2:1 tell', '0:1:1:2:0:0 tell', '0:1:1:2:1:0 elicit',
'0:1:1:2:1:1 else', '0:1:1:2:2:0 else', '0:1:1:2:2:1 elicit', '0:1:2:0:0:0 tell',
'0:1:2:0:0:1 tell', '0:1:2:0:1:0 elicit', '0:1:2:0:1:1 elicit', '0:1:2:0:2:1 elicit',
'0:1:2:1:0:0 tell', '0:1:2:1:0:1 tell', '0:1:2:1:1:0 elicit', '0:1:2:1:1:1 tell',
'0:1:2:1:2:0 tell', '0:1:2:1:2:1 elicit', '0:1:2:2:0:0 tell', '0:1:2:2:1:0 elicit',
'0:1:2:2:1:1 elicit', '0:1:2:2:2:0 tell', '0:1:2:2:2:1 elicit', '1:0:0:0:0:0 else',
'1:0:0:0:1:0 else', '1:0:0:0:1:1 else', '1:0:0:0:2:1 else', '1:0:0:1:0:0 else',
'1:0:0:1:1:0 elicit', '1:0:0:1:1:1 tell', '1:0:0:1:2:0 elicit', '1:0:0:1:2:1 else',
'1:0:0:2:0:0 else', '1:0:0:2:1:0 elicit', '1:0:0:2:1:1 elicit', '1:0:0:2:2:0 tell',

'1:0:0:2:2:1 elicit', '1:0:1:0:0:0 elicit', '1:0:1:0:1:0 tell', '1:0:1:0:1:1 elicit',
'1:0:1:0:2:1 elicit', '1:0:1:1:0:0 tell', '1:0:1:1:1:0 else', '1:0:1:1:1:1 elicit',
'1:0:1:1:2:1 tell', '1:0:1:2:0:0 elicit', '1:0:1:2:1:0 elicit', '1:0:1:2:2:0 elicit',
'1:0:1:2:2:1 tell', '1:0:2:0:0:0 tell', '1:0:2:0:1:0 tell', '1:0:2:0:1:1 tell',
'1:0:2:0:2:1 tell', '1:0:2:1:0:0 tell', '1:0:2:1:1:0 elicit', '1:0:2:1:1:1 tell',
'1:0:2:1:2:0 elicit', '1:0:2:1:2:1 elicit', '1:0:2:2:0:0 tell', '1:0:2:2:1:0 tell',
'1:0:2:2:1:1 elicit', '1:0:2:2:2:0 elicit', '1:0:2:2:2:1 elicit', '1:1:0:0:0:0 tell',
'1:1:0:0:0:1 else', '1:1:0:0:1:0 tell', '1:1:0:0:1:1 else', '1:1:0:0:2:0 elicit',
'1:1:0:0:2:1 elicit', '1:1:0:1:0:0 tell', '1:1:0:1:0:1 elicit', '1:1:0:1:1:0 else',
'1:1:0:1:1:1 tell', '1:1:0:1:2:0 elicit', '1:1:0:1:2:1 tell', '1:1:0:2:0:0 elicit',
'1:1:0:2:0:1 else', '1:1:0:2:1:0 tell', '1:1:0:2:1:1 else', '1:1:0:2:2:0 tell',
'1:1:0:2:2:1 elicit', '1:1:1:0:0:0 tell', '1:1:1:0:0:1 tell', '1:1:1:0:1:0 tell',
'1:1:1:0:1:1 elicit', '1:1:1:0:2:0 tell', '1:1:1:0:2:1 elicit', '1:1:1:1:0:0 elicit',
'1:1:1:1:0:1 elicit', '1:1:1:1:1:0 else', '1:1:1:1:1:1 tell', '1:1:1:1:2:0 tell',
'1:1:1:1:2:1 elicit', '1:1:1:2:0:0 tell', '1:1:1:2:0:1 elicit', '1:1:1:2:1:0 tell',
'1:1:1:2:1:1 tell', '1:1:1:2:2:0 tell', '1:1:1:2:2:1 tell', '1:1:2:0:0:0 elicit',
'1:1:2:0:1:0 elicit', '1:1:2:0:1:1 elicit', '1:1:2:0:2:1 tell', '1:1:2:1:0:0 tell',
'1:1:2:1:0:1 elicit', '1:1:2:1:1:0 elicit', '1:1:2:1:1:1 tell', '1:1:2:1:2:0 elicit',
'1:1:2:1:2:1 elicit', '1:1:2:2:0:0 tell', '1:1:2:2:1:0 tell', '1:1:2:2:1:1 tell',
'1:1:2:2:2:0 elicit']

Policy: '21skip'

Features: [18, 35, 45]

ECR: 46.744200

ECR Lower Bound: 41.828300

Mapping: [[0.30125, 0.6075], [0.25015, 0.45635], [0.25015, 0.45635]]

States: ['0:0:0 else', '0:1:1 else', '0:2:2 else', '1:0:0 else', '1:1:1 noskip', '1:2:2 else', '2:0:0 skip', '2:1:1 noskip', '2:2:2 noskip']

Policy: **'22et'**

Features: [24, 26, 27]

ECR: 79.291800

ECR Lower Bound: 58.21120

Mapping: [[0.067915], [21.303669], [2.2836]]

States: ['0:0:0 else', '0:0:1 elicit', '0:1:0 else', '0:1:1 elicit', '1:0:0 else', '1:0:1 elicit',
'1:1:0 else', '1:1:1 elicit']

Policy: **22skip'**

Features: [5, 19, 27, 32]

ECR: 95.363400

ECR Lower Bound: 88.012500

Mapping: [[631.9412], [0.25], [2.459, 6.635], [0.498302]]

States: ['0:0:0:0 noskip', '0:0:0:1 noskip', '0:0:1:0 noskip', '0:0:1:1 noskip', '0:0:2:0
noskip', '0:0:2:1 noskip', '1:0:0:0 noskip', '1:0:0:1 noskip', '1:0:1:0 noskip', '1:0:1:1
noskip', '1:0:2:0 noskip', '1:0:2:1 noskip', '1:1:0:0 noskip', '1:1:0:1 skip', '1:1:1:0
noskip', '1:1:1:1 noskip']

Policy: **'24et'**

Features: [14, 18]

ECR: 36.355700

ECR Lower Bound: 30.768000

Mapping: [[0.5], [0.52575]]

States: ['0:0 elicit', '0:1 tell', '1:0 elicit', '1:1 else']

Policy: **'24skip'**

Features: [17, 47]

ECR: 49.194000

ECR Lower Bound: 33.139500

Mapping: [[0.5], [13.31575]]

States: ['0:0 noskip', '0:1 noskip', '1:0 noskip', '1:1 skip']

Policy: **'27et'**

Features: [5, 9, 19, 27]

ECR: 37.696400

ECR Lower Bound: 28.755200

Mapping: [[76.037616], [3933.717329], [0.21775, 0.71775], [0.999999]]

States: ['0:0:0:0 else', '0:0:0:1 elicit', '0:0:1:0 else', '0:0:1:1 elicit', '0:0:2:0 elicit', '0:0:2:1 elicit', '0:1:0:0 elicit', '0:1:0:1 elicit', '0:1:1:0 else', '0:1:1:1 elicit', '0:1:2:0 tell', '0:1:2:1 elicit', '1:0:0:0 else', '1:0:0:1 elicit', '1:0:1:0 else', '1:0:1:1 elicit', '1:0:2:0 tell', '1:0:2:1 tell', '1:1:0:0 tell', '1:1:0:1 elicit', '1:1:1:0 else', '1:1:1:1 elicit', '1:1:2:0 tell', '1:1:2:1 tell']

Policy: **'27skip'**

Features: [9, 40, 42, 46, 56]

ECR: 60.041000

ECR Lower Bound: 40.637100

Mapping: [[4831.936247], [6.2033], [0.5205], [6.40425], [3.70513]]

States: ['0:0:0:0:0 skip', '0:0:0:0:1 skip', '0:0:0:1:0 noskip', '0:0:0:1:1 noskip', '0:0:1:0:0 skip', '0:0:1:0:1 skip', '0:1:0:1:0 noskip', '0:1:0:1:1 noskip', '0:1:1:0:0 noskip', '0:1:1:0:1 noskip', '0:1:1:1:1 noskip', '1:0:0:0:0 skip', '1:0:0:0:1 skip', '1:0:0:1:0 noskip', '1:0:0:1:1 skip', '1:0:1:0:0 noskip', '1:0:1:0:1 noskip', '1:1:0:1:0 noskip', '1:1:1:0:0 noskip', '1:1:1:0:1 noskip', '1:1:1:1:1 noskip']

Policy: '28et'

Features: [5, 18, 46]

ECR: 67.974200

ECR Lower Bound: 51.863700

Mapping: [[73.053098], [0.393393], [5.3352]]

States: ['0:0:0 else', '0:0:1 tell', '0:1:0 tell', '0:1:1 tell', '1:0:0 else', '1:0:1 tell', '1:1:0 elicit', '1:1:1 tell']

Policy: '28skip'

Features: [14, 15, 25, 46, 60]

ECR: 67.365600

ECR Lower Bound: 35.707600

Mapping: [[0.5], [15.10085], [22.754688, 23.282859], [5.191], [0.103777]]

States: ['0:0:0:0:0 noskip', '0:0:1:0:0 skip', '0:0:1:0:1 noskip', '0:1:0:0:0 noskip', '0:1:0:0:1 skip', '0:1:0:1:0 noskip', '0:1:0:1:1 noskip', '0:1:1:0:0 skip', '0:1:1:0:1 noskip', '0:1:1:1:0 skip', '0:1:1:1:1 noskip', '0:1:2:0:0 noskip', '0:1:2:0:1 skip', '0:1:2:1:0 noskip', '0:1:2:1:1 noskip', '1:0:0:0:0 skip', '1:0:0:0:1 noskip', '1:0:0:1:1 skip', '1:0:1:0:0 noskip', '1:0:1:0:1 noskip', '1:0:1:1:1 noskip', '1:0:2:0:0 skip', '1:0:2:0:1 noskip', '1:0:2:1:1 skip']

BIBLIOGRAPHY

- [Ai and Litman, 2009] Ai, H. and Litman, D. (2009). Setting up user action probabilities in user simulations for dialog system development. In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL)*. Suntec, Singapore.
- [Ainsworth and Fleming, 2005] Ainsworth, S. and Fleming, P. (2005). Evaluating a mixed-initiative authoring environment: Is redeem for real? In [Looi et al., 2005], pages 9–16.
- [Aleven et al., 2005] Aleven, V., McLaren, B. M., and Koedinger, K. R. (2005). Rapid development of computer-based tutors with the cognitive tutor authoring tools (ctat). In [Looi et al., 2005], page 990.
- [Aleven et al., 2006] Aleven, V., McLaren, B. M., Sewall, J., and Koedinger, K. R. (2006). The cognitive tutor authoring tools (ctat): Preliminary evaluation of efficiency gains. In [Ikeda et al., 2006], pages 61–70.
- [Aleven et al., 2004] Aleven, V., Ogan, A., Popescu, O., Torrey, C., and Koedinger, K. R. (2004). Evaluating the effectiveness of a tutorial dialogue system for self-explanation. In [Lester et al., 2004], pages 443–454.
- [Anderson et al., 1995] Anderson, J. R., Corbett, A. T., Koedinger, K. R., and Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4(2):167–207.
- [Aström, 1965] Aström, K. J. (1965). Optimal control of markov decision processes with the incomplete state estimation. *Journal of Computer and System Sciences*, 10:174–205.
- [Baker et al., 2004a] Baker, R. S., Corbett, A. T., and Koedinger, K. R. (2004a). Detecting student misuse of intelligent tutoring systems. In [Lester et al., 2004], pages 531–540.
- [Baker et al., 2004b] Baker, R. S., Corbett, A. T., Koedinger, K. R., and Wagner, A. Z. (2004b). Off-task behavior in the cognitive tutor classroom: when students "game the system". In Dykstra-Erickson, E. and Tscheligi, M., editors, *CHI*, pages 383–390. ACM.

- [Barnes and Stamper, 2008] Barnes, T. and Stamper, J. C. (2008). Toward automatic hint generation for logic proof tutoring using historical student data. In [Woolf et al., 2008], pages 373–382.
- [Beck et al., 2000] Beck, J., Woolf, B. P., and Beal, C. R. (2000). Advisor: A machine learning architecture for intelligent tutor construction. In *AAAI/IAAI*, pages 552–557. AAAI Press / The MIT Press.
- [Beck, 2001] Beck, J. E. (2001). *Advisor: A Machine-Learning Architecture for Intelligent Tutor Construction*. PhD thesis, Graduate School of the University of Massachusetts Amherst.
- [Beck and Mostow, 2008] Beck, J. E. and Mostow, J. (2008). How who should practice: Using learning decomposition to evaluate the efficacy of different types of practice for different types of students. In [Woolf et al., 2008], pages 353–362.
- [Bernsen and Dybkjaer, 1997] Bernsen, N. O. and Dybkjaer, L. (1997). *Designing Interactive Speech Systems: From First Ideas to User Testing*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- [Bloom, 1984] Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13:4–16.
- [Cade et al., 2008] Cade, W. L., Copeland, J. L., Person, N. K., and D’Mello, S. K. (2008). Dialogue modes in expert tutoring. In [Woolf et al., 2008], pages 470–479.
- [Calzolari et al., 2006] Calzolari, N., Cardie, C., and Isabelle, P., editors (2006). *ACL 2006, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, Sydney, Australia, 17-21 July 2006*. The Association for Computer Linguistics.
- [Cen et al., 2006] Cen, H., Koedinger, K. R., and Junker, B. (2006). Learning factors analysis - a general method for cognitive model evaluation and improvement. In [Ikeda et al., 2006], pages 164–175.
- [Cen et al., 2007] Cen, H., Koedinger, K. R., and Junker, B. (2007). Is over practice necessary? - improving learning efficiency with the cognitive tutor through educational data mining. In [Luckin et al., 2007], pages 511–518.
- [Chades et al., 2005] Chades, M. C., Garcia, F., and Sabbadin, R. (2005). Mdp toolbox v2.0 for matlab.
- [Chae et al., 2005] Chae, H. M., Kim, J. H., and Glass, M. (2005). Effective behaviors in a comparison between novice and expert algebra tutors. In *Proceedings of Sixteenth Midwest AI and Cognitive Science Conference, Dayton*, pages 25–30.

- [Chi et al., 2008a] Chi, M., Jordan, P. W., VanLehn, K., and Hall, M. (2008a). Reinforcement learning-based feature selection for developing pedagogically effective tutorial dialogue tactics. In de Baker, R. S. J., Barnes, T., and Beck, J. E., editors, *EDM*, pages 258–265. www.educationaldatamining.org.
- [Chi and VanLehn, 2008] Chi, M. and VanLehn, K. (2008). Eliminating the gap between the high and low students through meta-cognitive strategy instruction. In [Woolf et al., 2008], pages 603–613.
- [Chi et al., 1994] Chi, M. T. H., de Leeuw, N., Chiu, M.-H., and LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18(3):439–477.
- [Chi et al., 2008b] Chi, M. T. H., Roy, M., and Hausmann, R. G. M. (2008b). Observing tutorial dialogues collaboratively: Insights about human tutoring effectiveness from vicarious learning. *Cognitive Science*, 32(2):301–342.
- [Chi et al., 2004] Chi, M. T. H., Siler, S., and Jeong, H. (2004). Can tutors monitor students’ understanding accurately? *Cognition and Instruction*, 22(3):363–387.
- [Chi et al., 2001] Chi, M. T. H., Siler, S., Jeong, H., Yamauchi, T., and Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science*, 25:471–533.
- [Cho et al., 2000] Cho, B.-I., Michael, J. A., Rovick, A. A., and Evens, M. W. (2000). An analysis of multiple tutoring protocols. In Gauthier, G., Frasson, C., and VanLehn, K., editors, *Intelligent Tutoring Systems*, volume 1839 of *Lecture Notes in Computer Science*, pages 212–221. Springer.
- [Clark et al., 1976] Clark, C. M., Snow, R. E., and Shavelson, R. J. (1976). Three experiments on learning to teach. *Journal of Teacher Education*, 27:174–180.
- [Cohen et al., 1982] Cohen, P. A., Kulik, J. A., and Kulik, C.-L. C. (1982). Educational outcomes of tutoring: A meta-analysis of findings. *American Educational Research Journal*, 19(2):237–248.
- [Coley, 2001] Coley, R. J. (2001). Differences in the gender gap: Comparisons across racial/ethnic groups in education and work. policy information report. Technical report, Educational Testing Service, Rosedale Road, Princeton, NJ 08541-0001 609-734-5694; (www.ets.org/research).
- [Collins et al., 1989] Collins, A., Brown, J. S., and Newman, S. E. (1989). Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics. In Resnick, L. B., editor, *Knowing, learning and instruction: Essays in honor of Robert Glaser*, chapter 14, pages 453–494. Lawrence Erlbaum Associates: Hillsdale New Jersey.
- [Collins and Stevens, 1982] Collins, A. and Stevens, A. (1982). Goals and strategies for inquiry teachers. *Advances in Instructional Psychology*, 2:65–119.

- [Conati and VanLehn, 2000] Conati, C. and VanLehn, K. (2000). Toward computer-based support and cristina conati and kurt vanlehn. *International Journal of Artificial Intelligence in Education*, 11:398–415.
- [Core et al., 2003] Core, M. G., Moore, J. D., and Zinn, C. (2003). The role of initiative in tutorial dialogue. In *EACL*, pages 67–74.
- [Eugenio et al., 2006] Eugenio, B. D., Kershaw, T. C., Lu, X., Corrigan-Halpern, A., and Ohlsson, S. (2006). Toward a computational model of expert tutoring: A first report. In Sutcliffe, G. and Goebel, R., editors, *FLAIRS Conference*, pages 503–508. AAAI Press.
- [Evens and Michael, 2006] Evens, M. and Michael, J. (2006). *One-on-one Tutoring By Humans and Machines*. Mahwah, NJ: Erlbaum.
- [Forbes-Riley et al., 2007] Forbes-Riley, K., Litman, D. J., Purandare, A., Rotaru, M., and Tetreault, J. R. (2007). Comparing linguistic features for modeling learning in computer tutoring. In [Luckin et al., 2007], pages 270–277.
- [Frampton and Lemon, 2005] Frampton, M. and Lemon, O. (2005). Reinforcement learning of dialogue strategies using the user’s last dialogue act. In *Proceedings of the IJCAI Workshop on K&R in Practical Dialogue Systems*, pages 62–67.
- [Frampton and Lemon, 2006] Frampton, M. and Lemon, O. (2006). Learning more effective dialogue strategies using limited dialogue move features. In [Calzolari et al., 2006].
- [Gallagher, 2001] Gallagher, T. (2001). Equal opportunities commission conference on boys and girls in the 21st century: Gender differences in learning. Technical report, Equal Opportunities Commission.
- [Graesser et al., 1995] Graesser, A. C., Person, N., and Magliano, J. (1995). Collaborative dialog patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychology*, 9:359–387.
- [Graesser et al., 2001] Graesser, A. C., VanLehn, K., Rosé, C. P., Jordan, P. W., and Harter, D. (2001). Intelligent tutoring systems with conversational dialogue. *AI Magazine*, 22(4):39–52.
- [Hauskrecht, 1997] Hauskrecht, M. (1997). *Planning and control in stochastic domains with imperfect information*. PhD thesis, MIT. Available as Technical Report: MIT-LCS-TR-738, 1997.
- [Henderson et al., 2005] Henderson, J., Lemon, O., and Georgila, K. (2005). Hybrid reinforcement/supervised learning for dialogue policies from communicator data. In *IJCAI Workshop on K&R in Practical Dialogue Systems*, pages 68–75.
- [Hume et al., 1995] Hume, G., Michael, J., Rovick, A., and Evens, M. (1995). Controlling active learning: how tutors decide when to generate hints. In *Proceedings of the 8th Florida Artificial Intelligence Research Symposium.*, pages 157–161.

- [Ikeda et al., 2006] Ikeda, M., Ashley, K. D., and Chan, T.-W., editors (2006). *Intelligent Tutoring Systems, 8th International Conference, ITS 2006, Jhongli, Taiwan, June 26-30, 2006, Proceedings*, volume 4053 of *Lecture Notes in Computer Science*. Springer.
- [Janarthanam and Lemon, 2009] Janarthanam, S. and Lemon, O. (2009). User simulations for online adaptation and knowledge-alignment in troubleshooting dialogue systems. In *Proceedings of the 12th SEMdial Workshop on on the Semantics and Pragmatics of Dialogues*.
- [Jolliffe, 2002] Jolliffe, I. T. (2002). *Principal Component Analysis, Series: Springer Series in Statistics*. Springer, New York, 2nd edition.
- [Jordan et al., 2007] Jordan, P. W., Hall, B., Ringenberg, M. A., Cue, Y., and Rosé, C. P. (2007). Tools for authoring a dialogue agent that participates in learning studies. In [Luckin et al., 2007], pages 43–50.
- [Jordan et al., 2006] Jordan, P. W., Ringenberg, M. A., and Hall, B. (2006). Tools for authoring a dialogue agent that participates in learning studies. In *Proceedings of ITS06 Workshop on Teaching with Robots, Agents, and NLP*.
- [Katz et al., 2007] Katz, S., Connelly, J., and Wilson, C. (2007). Out of the lab and into the classroom: An evaluation of reflective dialogue in andes. In [Luckin et al., 2007], pages 425–432.
- [Katz et al., 2000] Katz, S., O’Donnell, G., and Kay, H. (2000). An approach to analyzing the role and structure of reflective dialogue. *International Journal of Artificial Intelligence and Education*, 11:320–343.
- [Kim et al., 2005] Kim, J. H., Chae, H. M., and Glass, M. (2005). Expert and novice algebra tutor behaviors compared (poster abstract). In *Proceedings of the 27th Annual Conference of the Cognitive Science Society, COGSCI 2005, Stresa, Italy*.
- [Koedinger and Alevan, 2007] Koedinger, K. R. and Alevan, V. (2007). Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review*, 19(3):239–264.
- [Koedinger et al., 1997] Koedinger, K. R., Anderson, J. R., Hadley, W. H., and Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8(1):30–43.
- [Lane and VanLehn, 2005] Lane, H. C. and VanLehn, K. (2005). Teaching the tacit knowledge of programming to novices with natural language tutoring. *Computer Science Education*, 15(3):183–201.
- [Lemon et al., 2006] Lemon, O., Georgila, K., and Henderson, J. (2006). Evaluating effectiveness and portability of reinforcement learned dialogue strategies with real users: the talk towninfo evaluation. In *IEEE/ACL Spoken Language Technology*.

- [Lester et al., 2004] Lester, J. C., Vicari, R. M., and Paraguaçu, F., editors (2004). *Intelligent Tutoring Systems, 7th International Conference, ITS 2004, Maceió, Alagoas, Brazil, August 30 - September 3, 2004, Proceedings*, volume 3220 of *Lecture Notes in Computer Science*. Springer.
- [Levin and Pieraccini, 1997] Levin, E. and Pieraccini, R. (1997). A stochastic model of computer-human interaction for learning dialogue strategies. In *In EUROSPEECH 97*, pages 1883–1886.
- [Litman and Silliman, 2004] Litman, D. J. and Silliman, S. (2004). Itspoke: an intelligent tutoring spoken dialogue system. In *HLT-NAACL '04: Demonstration Papers at HLT-NAACL 2004 on XX*, pages 5–8, Morristown, NJ, USA. Association for Computational Linguistics.
- [Looi et al., 2005] Looi, C.-K., McCalla, G. I., Bredeweg, B., and Breuker, J., editors (2005). *Artificial Intelligence in Education - Supporting Learning through Intelligent and Socially Informed Technology, Proceedings of the 12th International Conference on Artificial Intelligence in Education, AIED 2005, July 18-22, 2005, Amsterdam, The Netherlands*, volume 125 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- [Lu et al., 2007] Lu, X., Eugenio, B. D., Kershaw, T. C., Ohlsson, S., and Corrigan-Halpern, A. (2007). Expert vs. non-expert tutoring: Dialogue moves, interaction patterns and multi-utterance turns. In Gelbukh, A. F., editor, *CICLing*, volume 4394 of *Lecture Notes in Computer Science*, pages 456–467. Springer.
- [Luckin et al., 2007] Luckin, R., Koedinger, K. R., and Greer, J. E., editors (2007). *Artificial Intelligence in Education, Building Technology Rich Learning Contexts That Work, Proceedings of the 13th International Conference on Artificial Intelligence in Education, AIED 2007, July 9-13, 2007, Los Angeles, California, USA*, volume 158 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- [McArthur et al., 1982] McArthur, D., Stasz, C., and Zmuidzinas, M. (1982). Tutoring techniques in algebra. *Cognition and Instruction*, 7(3):197–244.
- [Merrill et al., 1995] Merrill, D. C., Reiser, B. J., Merrill, S. K., and Landes, S. (1995). Tutoring: Guided learning by doing. *Cognition and Instruction*, 13(3):315–372.
- [Merrill et al., 1992] Merrill, D. C., Reiser, B. J., Ranney, M., and Trafton, J. G. (1992). Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences*, 2(3):277–306.
- [Moore et al., 2004] Moore, J. D., Porayska-Pomsta, K., Varges, S., and Zinn, C. (2004). Generating tutorial feedback with affect. In Barr, V. and Markov, Z., editors, *FLAIRS Conference*. AAAI Press.

- [Murray and VanLehn, 2006] Murray, R. C. and VanLehn, K. (2006). A comparison of decision-theoretic, fixed-policy and random tutorial action selection. In [Ikeda et al., 2006], pages 114–123.
- [Murray et al., 2003] Murray, T., Blessing, S., and Ainsworth, S. (2003). *Authoring Tools for Advanced Technology Learning Environments*. Kluwer Academic/Springer Pub.: Netherlands.
- [Noe, 1986] Noe, R. A. (1986). Trainee’s attributes and attitudes: neglected influences on training effectiveness’. *Academy of Management Review*, 11:736–749.
- [Paek and Chickering, 2005] Paek, T. and Chickering, D. (2005). The markov assumption in spoken dialogue management. In *6th SIGDial Workshop on Discourse and Dialogue*.
- [Purandare and Litman, 2008] Purandare, A. and Litman, D. J. (2008). Content-learning correlations in spoken tutoring dialogs at word, turn, and discourse levels. In Wilson, D. and Lane, H. C., editors, *FLAIRS Conference*, pages 439–443. AAAI Press.
- [Putnam, 1987] Putnam, R. T. (1987). Structuring and adjusting content for students: A study of live and simulated tutoring of addition. *American Educational Research Journal*, 24(1):13–48.
- [Quek et al., 2002] Quek, C. L., Wong, A. F., and Fraser, B. J. (2002). Gender differences in the perceptions of chemistry laboratory classroom environments. *Queensland Journal of Educational Research*, 18:164–182.
- [Raux et al., 2005] Raux, A., , Langner, Bohus, D., and Eskenazi, M. (2005). Let’s go public! taking a spoken dialog system to the real world. In *Proceedings of Interspeech (Eurospeech)*.
- [Reif and Scott, 1999] Reif, F. and Scott, L. A. (1999). Teaching scientific thinking skills: Students and computers coaching each other. *American Journal of Physics*, 67(9):819–831.
- [Rieser and Lemon, 2006] Rieser, V. and Lemon, O. (2006). Using machine learning to explore human multimodal clarification strategies. In [Calzolari et al., 2006].
- [Rose et al., 2001] Rose, C. P., Moore, J. D., VanLehn, K., and Allbritton, D. (2001). A comparative evaluation of socratic versus didactic tutoring. In *Proceedings of Cognitive Sciences Society*, pages 869–874.
- [Rudnicky et al., 1999] Rudnicky, A., Thayer, E., Constantinides, P., Tchou, C., Shern, R., Lenzo, K., Xu, W., and Oh, A. (1999). Creating natural dialogs in the carnegie mellon communicator system. In *Proceedings of Eurospeech.*, volume 4, pages 1531–1534.
- [Singh et al., 1999] Singh, S. P., Kearns, M. J., Litman, D. J., and Walker, M. A. (1999). Reinforcement learning for spoken dialogue systems. In Solla, S. A., Leen, T. K., and Müller, K.-R., editors, *NIPS*, pages 956–962. The MIT Press.

- [Singh et al., 2002] Singh, S. P., Litman, D. J., Kearns, M. J., and Walker, M. A. (2002). Optimizing dialogue management with reinforcement learning: Experiments with the njfun system. *J. Artif. Intell. Res. (JAIR)*, 16:105–133.
- [Sleeman et al., 1989] Sleeman, D. H., Kelly, A. E., Martinak, R., Ward, R. D., and Moore, J. L. (1989). Studies of diagnosis and remediation with high school algebra students. *Cognitive Science*, 13(4):551–568.
- [Stamper et al., 2007] Stamper, J. C., Barnes, T., and Croy, M. J. (2007). Extracting student models for intelligent tutoring systems. In *AAAI*, pages 1900–1901. AAAI Press.
- [Sutton and Barto, 1998] Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning*. MIT Press Bradford Books.
- [Swerts et al., 2000] Swerts, M., Litman, D., and Hirshberg, J. (2000). Corrections in spoken dialogue systems. In *Proceedings of the Sixth International Conference on Spoken Language Processing (ICSLP 2000)*, volume 2, pages 615–618.
- [Tetreault and Litman, 2006a] Tetreault, J. and Litman, D. (2006a). Using reinforcement learning to build a better model of dialogue state. In *Proceedings 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL), Trento, Italy*.
- [Tetreault et al., 2007] Tetreault, J. R., Bohus, D., and Litman, D. J. (2007). Estimating the reliability of mdp policies: a confidence interval approach. In Sidner, C. L., Schultz, T., Stone, M., and Zhai, C., editors, *HLT-NAACL*, pages 276–283. The Association for Computational Linguistics.
- [Tetreault and Litman, 2006b] Tetreault, J. R. and Litman, D. J. (2006b). Comparing the utility of state features in spoken dialogue using reinforcement learning. In Moore, R. C., Bilmes, J. A., Chu-Carroll, J., and Sanderson, M., editors, *HLT-NAACL*. The Association for Computational Linguistics.
- [Tetreault and Litman, 2008] Tetreault, J. R. and Litman, D. J. (2008). A reinforcement learning approach to evaluating state representations in spoken dialogue systems. *Speech Communication*, 50(8-9):683–696.
- [VanLehn, 1999] VanLehn, K. (1999). Rule learning events in the acquisition of a complex skill: An evaluation of cascade. *Journal of the Learning Sciences*, 8(1):71–125.
- [VanLehn, 2006] VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, 16(3):227–265.
- [VanLehn, 2009] VanLehn, K. (2009). The two-sigma effect revisited: A meta-analysis of human tutoring and several types of computer tutoring. (submitted).

- [VanLehn et al., 2007a] VanLehn, K., Graesser, A. C., Jackson, G. T., Jordan, P., Olney, A., and Rose, C. P. (2007a). When are tutorial dialogues more effective than reading? *Cognitive Science*, 31(1):3–62.
- [VanLehn et al., 2007b] VanLehn, K., Jordan, P., and Litman, D. (2007b). Developing pedagogically effective tutorial dialogue tactics: Experiments and a testbed. In *Proceedings of SLATE Workshop on Speech and Language Technology in Education ISCA Tutorial and Research Workshop*.
- [VanLehn et al., 2002] VanLehn, K., Jordan, P. W., Rosé, C. P., Bhembe, D., Böttner, M., Gaydos, A., Makatchev, M., Pappuswamy, U., Ringenberg, M. A., Roque, A., Siler, S., and Srivastava, R. (2002). The architecture of why2-atlas: A coach for qualitative physics essay writing. In Cerri, S. A., Gouardères, G., and Paraguaçu, F., editors, *Intelligent Tutoring Systems*, volume 2363 of *Lecture Notes in Computer Science*, pages 158–167. Springer.
- [VanLehn et al., 2005] VanLehn, K., Lynch, C., Schulze, K., Shapiro, J. A., Shelby, R., Taylor, L., Treacy, D., Weinstein, A., and Wintersgill, M. (2005). The andes physics tutoring system: Lessons learned. *Int. J. Artif. Intell. Ed.*, 15(3):147–204.
- [VanLehn et al., 2003] VanLehn, K., Siler, S., Murray, R. C., Yamauchi, T., and Baggett, W. B. (2003). Why do only some events cause learning during human tutoring? *Cognition and Instruction*, 21(3):209–249.
- [Vygotsky, 1971] Vygotsky, L. (1971). Interaction between learning and development. In Cole, T. M., editor, *In Mind in Society.*, pages 79–91. Harvard University Press: Cambridge Massachusetts.
- [Walker, 2000] Walker, M. A. (2000). An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research*, 12:387–416.
- [Williams et al., 2005] Williams, J., Poupart, P., and Young, S. (2005). Factored partially observable markov decision processes for dialogue management. In *Int. Joint Conf. on Artificial Intelligence*.
- [Williams and Young, 2007a] Williams, J. and Young, S. (2007a). Partially observable markov decision processes for spoken dialog systems. *Computer Speech and Language*, 21(2):231–422.
- [Williams and Young, 2007b] Williams, J. and Young, S. (2007b). Scaling pomdps for spoken dialog management. *IEEE Trans. on Audio, Speech, and Language Processing*.
- [Woolf et al., 2008] Woolf, B. P., Aïmeur, E., Nkambou, R., and Lajoie, S. P., editors (2008). *Intelligent Tutoring Systems, 9th International Conference, ITS 2008, Montreal, Canada, June 23-27, 2008, Proceedings*, volume 5091 of *Lecture Notes in Computer Science*. Springer.

[Young, 1999] Young, S. (1999). Probabilistic methods in spoken dialogue systems. *Philosophical Transactions of the Royal Society (Series A)*, 358:1389–1402.