

## ABSTRACT

ROWE, JONATHAN PAUL. Narrative-Centered Tutorial Planning with Concurrent Markov Decision Processes. (Under the direction of Dr. James C. Lester.)

Narrative-centered learning environments are a class of educational games that embed learning in interactive story scenarios. A key promise of narrative-centered learning environments is the ability to tightly integrate educational subject matter with engaging stories. While narrative-centered learning environments offer motivational benefits, students may struggle to employ effective problem-solving strategies in these settings without tailored pedagogical guidance. A potential solution is devising narrative-centered tutorial planners, which dynamically generate and revise story events in accordance with educational goals and students' needs. By augmenting interactive narratives to provide targeted remediation and story-embedded scaffolding, narrative-centered tutorial planners can deliver individualized pedagogical guidance while preserving the motivational benefits of narrative-centered learning.

This dissertation presents a data-driven framework for devising narrative-centered tutorial planners with concurrent Markov decision processes. The framework decomposes interactive narratives into several independent *adaptable event sequences (AESs)*, which are modularized story structures that can be realized in multiple forms within an interactive narrative. Dynamically adapting an event sequence is modeled as a Markov decision process (MDP), with rewards computed from students' experiential outcomes, such as learning gains and problem-solving results. Distinct policies are independently machine-learned for each AES, and the policies are then re-combined using arbitration procedures to form a composite narrative-centered tutorial planner. The proposed framework is well suited to machine learning interactive narrative planners directly from actual students' data. By inducing models from student data, the approach offers promise for yielding effective and practical models, as well as insights into student learning and design of narrative-centered learning environments.

We investigate the framework's effectiveness by examining an implemented narrative-centered tutorial planner for the CRYSTAL ISLAND learning environment. In order to

induce the planning model, a sizable training corpus was generated through a series of studies involving middle school students interacting with a version of CRYSTAL ISLAND imbued with an exploratory narrative-centered tutorial planner. Narrative adaptation policies for each AES were induced using certainty equivalent learning. Afterward, we conducted a controlled experiment with a group of middle school students to evaluate the induced narrative-centered tutorial planner's effectiveness at enhancing students' problem-solving processes. The experiment revealed that the Induced Planner significantly improved students' efficiency and deliberateness in problem solving compared to a baseline system. Specifically, students solved the mystery in less time when interacting with the induced planner, and they demonstrated improved information gathering and hypothesis testing behaviors. The results provide evidence in support of our thesis that modeling narrative-centered tutorial planning with concurrent Markov decision processes is an effective approach for dynamically tailoring story events to enhance student problem solving in narrative-centered learning environments.

Narrative-Centered Tutorial Planning with Concurrent Markov Decision Processes

by  
Jonathan Paul Rowe

A dissertation submitted to the Graduate Faculty of  
North Carolina State University  
in partial fulfillment of the  
requirements for the Degree of  
Doctor of Philosophy

Computer Science

Raleigh, North Carolina

2013

APPROVED BY:

---

Dr. James C. Lester  
Committee Chair

---

Dr. R. Michael Young

---

Dr. David L. Roberts

---

Mr. Patrick J. Fitzgerald

## **DEDICATION**

To my parents, who provided the foundation for all I do,  
my wife Kristen, whose love and support make it possible,  
and my daughter Azalea, who is the reason for doing it.

## BIOGRAPHY

Jonathan Paul Rowe was born in Reading, Pennsylvania on August 5, 1984. He received his first computer, a Compaq PC with an 80386 Intel processor, from his stepfather, a refurbished hand-me-down from the offices at Connors Investor Services. During these years, Jonathan developed a passion for computers and videogames, deciding that he would one day make videogames for a living. In addition to these interests, Jonathan was active in the Boy Scouts of America, eventually earning its highest award: the Eagle Scout rank.

He graduated from Schuylkill Valley High School in 2002, and went on to attend Lafayette College in Easton, Pennsylvania. At Lafayette, he majored in Computer Science with a minor in Mathematics. In addition to his classwork, Jonathan was a member of the College Men's Track and Field team, where he served as the Men's Team Captain for two years, and was eventually named to ESPN The Magazine Academic All-District II First Team. Jonathan earned his Bachelor of Science degree from Lafayette in 2006, graduating summa cum laude.

The next Fall, Jonathan began his graduate career at North Carolina State University. During his time in graduate school, he had two stints at summer internships: in 2007 he interned at USC's Institute for Creative Technologies researching story-authoring support tools, and in 2011 he interned at the SAS Institute's Education Practice working on mobile learning applications. He earned his Master of Science degree in Computer Science in 2010.

Jonathan married his wife Kristen in fall 2009. Their first daughter, Azalea Lynn, was born in 2012. As planned, Jonathan now makes videogames for a living.

## ACKNOWLEDGMENTS

First, I would like to thank my advisor James Lester. James has provided an exceptional model of scholarly professionalism throughout my graduate career. Further, his collegiality, enthusiasm, and friendship have left profound imprints that extend far beyond this dissertation. I would also like to thank the members of my committee, Michael Young, Dave Roberts, and Pat Fitzgerald, for their encouragement, insights, and feedback. They have each influenced and strengthened this dissertation considerably. Thanks as well to Rob St. Amant, who helped to inspire my interests in human-computer interaction, and provided crucial assistance during the final stages of my dissertation process.

I have been privileged to work with several great mentors throughout my graduate career, and in each case their work has deeply influenced my own. I am particularly indebted to Brad Mott, whose insights, encouragement, and technical expertise have proven indispensable throughout the research described in this dissertation. Scott McQuiggan and Mark Riedl also provided valuable mentorship during research opportunities outside the brick-laden walls of NCSU, substantially broadening my view of scholarly research.

It has been a pleasure to work with so many bright colleagues in the IntelliMedia Group over the years, including Alok Baikadi, Kristy Boyer, Stephen Cossa, Kirby Culbertson, Julius Goth, Joe Grafsgaard, Eunyoung Ha, Sarah Hegler, Seung Lee, Karoon McDowell, Chris Mitchell, Sam Leeman-Munk, Eleni Lobene, Wookhee Min, Adam Osgood, Justin Phillips, Marc Russo, Jennifer Sabourin, Lucy Shores, Rob Taylor, Andy Smith, Adam Winsor, and Donnie Wrights. I have been fortunate to co-author papers and collaborate on projects with nearly all of them, and it has been through these processes that most of my learning in graduate school has occurred. Special thanks go to Rob Taylor for his late-night contributions to the software described in this dissertation. In addition, many thanks to other friends and colleagues from NCSU, including Travis Breaux, Dave Crist,

Arnav Jhala, Alex Kuhl, Jim Thomas, and Stephen Ware for their support, camaraderie, and stimulating conversations throughout the many stages of my graduate career.

I have been fortunate to participate in collaborative research with several faculty members from other colleges and departments at NCSU. In particular, I am grateful to Hiller Spires, who made the initial connections with schools interested in using CRYSTAL ISLAND in their classrooms. I also wish to thank John Nietfeld and James Minogue for many stimulating conversations about student learning, instruction, and life.

This research would not have been possible without the teachers, students, and administrators of Johnston County Schools and Centennial Campus Magnet Middle School. In particular, I wish to thank Elwood Peters, Chris Godwin, Vickie Boyette, Kathy Feola, Rebecah Spraggon, Kelly Mance, Heidi Adams, Michelle Zimmerman, and Divern Combs.

Many thanks to the IT Staff at the William & Ida Friday Institute, particularly Bob Haggerty and Mustapha Jawara, for their support during the multiple studies with CRYSTAL ISLAND reported in this dissertation. Furthermore, thanks to the staff of the Department of Computer Science, with special thanks to Carol Allen, Susan Peaslee, Linda Honeycutt, and Margery Page for their frequent assistance involving matters large and small.

Finally, I want to thank my wife Kristen, whose love, patience, and support have remained steadfast throughout the many late nights and years of graduate school. Throughout my graduate career, she has selflessly taken on extra household responsibilities, while juggling her own professional career, to support me as I have pursued a Ph.D. I also wish to thank my parents, Sherry and Ron Seher, for their regular encouragement and unconditional love. Without their sacrifices and influence, I would not be where I am today.

The work reported in this dissertation was supported by the National Science Foundation under grants REC-0632450, IIS-0812291, and DRL-0822200. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

## TABLE OF CONTENTS

<b>LIST OF TABLES .....</b>	<b>viii</b>
<b>LIST OF FIGURES .....</b>	<b>ix</b>
<b>CHAPTER 1 Introduction .....</b>	<b>1</b>
1.1 Problem .....	2
1.2 Approach.....	3
1.3 Thesis Statement and Hypotheses.....	6
1.4 Contributions.....	7
1.5 Organization .....	9
<b>CHAPTER 2 Background and Related Work .....</b>	<b>10</b>
2.1 Narrative-Centered Learning Environments.....	10
2.2 Intelligent Tutoring Systems .....	19
2.3 Dynamic Adaptation of Interactive Narratives.....	22
2.4 Decision-Theoretic Approaches for Narrative Adaptation .....	25
<b>CHAPTER 3 Inducing Narrative-Centered Tutorial Planners.....</b>	<b>30</b>
3.1 Modular Reinforcement Learning.....	30
3.2 Narrative-Centered Tutorial Planning with Concurrent Markov Decision Processes .....	36
3.3 Inducing Narrative-Centered Tutorial Planning Policies .....	52
3.4 Theoretical Analysis.....	57
<b>CHAPTER 4 CRYSTAL ISLAND .....</b>	<b>63</b>
4.1 CRYSTAL ISLAND Narrative-Centered Learning Environment .....	64
4.2 Human Subject Study .....	70
4.3 Student Learning, Problem-Solving and Engagement Findings .....	74
4.4 Individual Difference Findings .....	77
<b>CHAPTER 5 Training Corpus Collection.....</b>	<b>83</b>
5.1 Modified CRYSTAL ISLAND.....	84
5.2 Corpus Collection Studies .....	90
5.2.1 Population.....	91
5.2.2 Design.....	91
5.2.3 Pre-Study Materials.....	92
5.2.4 Post-Study Materials .....	93
5.2.5 Apparatus .....	94
5.2.6 Procedure.....	95
5.3 Structure of Narrative-Centered Tutorial Planning Corpus .....	96
<b>CHAPTER 6 Evaluation Study .....</b>	<b>101</b>
6.1 An Implemented Narrative-Centered Tutorial Planner .....	101
6.2 Evaluation Experiment.....	108
6.2.1 Population.....	109
6.2.2 Design.....	109

6.2.3 Procedure.....	110
6.3 Profile of Induced Narrative-Centered Tutorial Planner’s Decisions .....	111
<b>CHAPTER 7 Results .....</b>	<b>118</b>
7.1 Content Learning Results.....	118
7.2 Problem Solving Results .....	119
7.3 Engagement Results .....	129
<b>CHAPTER 8 Discussion .....</b>	<b>132</b>
8.1 Enhancing Students’ Problem-Solving Processes with Narrative Adaptations .....	133
8.2 Limitations .....	138
<b>CHAPTER 9 Conclusion .....</b>	<b>142</b>
9.1 Hypotheses Revisited .....	143
9.2 Summary .....	145
9.3 Future Work .....	147
9.4 Concluding Remarks .....	149
<b>REFERENCES.....</b>	<b>151</b>
<b>APPENDICES .....</b>	<b>167</b>
Appendix A: Demographics and Game-Playing Experience Survey .....	168
Appendix B: Self-Efficacy for Science Questionnaire .....	170
Appendix C: CRYSTAL ISLAND Curriculum Test v2.0 .....	171
Appendix D: Adapted Perceived Interest Questionnaire .....	175
Appendix E: CRYSTAL ISLAND Handout Materials.....	176
Appendix F: CRYSTAL ISLAND Curriculum Test v3.0 .....	179
Appendix G: CRYSTAL ISLAND Narrative Questionnaire .....	184
Appendix H: Induced Narrative Adaptation Policies.....	188

## LIST OF TABLES

Table 1. Procedure for modeling a narrative-centered tutorial planner with concurrent Markov decision processes .....	62
Table 2. Point values for calculating final game score .....	73
Table 3. Regression predicting students' microbiology post-test scores .....	75
Table 4. Hierarchical regression predicting students' microbiology post-test scores .....	76
Table 5. Breakdown of student actions in the planning corpus .....	97
Table 6. Breakdown of narrative adaptations in the planning corpus.....	98
Table 7. Contingency table for Quentin's Revelation AES .....	113
Table 8. Contingency table for Bryce's Revelation AES.....	113
Table 9. Contingency table for Mystery Solution AES .....	114
Table 10. Contingency table for Initial Lab Test Count AES.....	114
Table 11. Breakdown of narrative adaptation frequencies for recurring AESs.....	116

## LIST OF FIGURES

Figure 1. Student interaction with a narrative-centered tutorial planner .....	4
Figure 2. Agent-environment architecture for classical reinforcement learning .....	31
Figure 3. Agent-environment architecture for narrative-centered tutorial planner based on modular reinforcement learning.....	38
Figure 4. Feature selection for narrative-centered tutorial planner states.....	46
Figure 5. Partial order among sub-goals in the CRYSTAL ISLAND narrative scenario.....	65
Figure 6. Conversation with CRYSTAL ISLAND's camp nurse .....	66
Figure 7. Diagnosis worksheet for recording findings in CRYSTAL ISLAND.....	67
Figure 8. Interior of the camp infirmary in CRYSTAL ISLAND .....	69
Figure 9. Testing device in the camp's laboratory.....	69
Figure 10. Distribution of students' final game scores .....	78
Figure 11. Individual differences prior to gameplay, in-game, and after gameplay.....	79
Figure 12. State representation for induced narrative-centered tutorial planner .....	103
Figure 13. Students' conversational behaviors with the virus specialist .....	122
Figure 14. Students' conversational behaviors with the bacteria specialist.....	122
Figure 15. Students' conversational behaviors with the camp nurse.....	123
Figure 16. Students' conversational behaviors with the sick lead scientist.....	124
Figure 17. Students' conversational behaviors with the camp cook.....	124
Figure 18. Students' Salmonellosis poster-viewing behaviors .....	125
Figure 19. Students' Anthrax poster-viewing behaviors .....	126
Figure 20. Students' Botulism poster-viewing behaviors.....	126
Figure 21. Students' Ebola poster-viewing behaviors .....	127

# CHAPTER 1

## Introduction

The past decade has witnessed a growing recognition of the potential for digital games to deliver effective and engaging learning experiences. In 2006, experts convened at the National Summit on Educational Games agreed “gaming could help address one of the nation’s most pressing needs — strengthening our system of education and preparing workers for 21<sup>st</sup> century jobs” (Federation of American Scientists, 2006). The National Academies’ Board of Science Education recently released a report that asserted “[Games] enable learners to see and interact with representations of natural phenomena that would otherwise be impossible to observe ... games can motivate learners with challenges and rapid feedback and tailor instruction to individual learners’ needs and interests” (Honey & Hilton, 2011). In addition to commissioned reports, there is a growing literature discussing the pedagogical potential of educational games (Gee, 2007; Prensky, 2001; Shaffer, 2006). This literature argues that well-designed games inherently promote learning, and the skills taught by games are vital for success in the modern economy, such as critical thinking, communication, teamwork, and problem solving.

One particularly promising class of educational games is narrative-centered learning environments. Narrative-centered learning environments embed educational content and problem solving with interactive story scenarios. They combine salient features of stories (rich settings, believable characters, and compelling plots) with key elements of digital games (agency, rewards, and multimedia feedback), offering significant promise for increasing student motivation, supporting meaning making, and guiding complex problem solving. In narrative-centered learning environments, students become active participants in ongoing narratives, acquiring and applying knowledge as they advance interactive plots. Narrative-centered learning environments enable students to explore rich virtual

environments, converse with casts of believable characters, and experience fantasy events that cannot be easily emulated in classrooms, all in the service of promoting effective and engaging learning experiences.

In order to promote learning, narrative-centered learning environments tap into students' innate facilities for crafting and understanding stories (Bruner, 1991; Graesser & Ottati, 1996; Polkinghorne, 1988). Narrative-centered learning environments provide meaningful contexts for problem-solving activities, which illustrate connections between theories and applications (Jonassen & Hernandez-Serrano, 2002). In addition, narrative-centered learning environments encourage engaged learning by tightly integrating pedagogy and narrative elements. In recent years, narrative-centered learning environments have been investigated for a range of domains, such as language learning (Johnson, 2010), anti-bullying education (Vannini, et al., 2011), intercultural negotiation training (Kim, et al., 2009), middle school science (Ketelhut, Dede, Clarke, & Nelson, 2010), and network security (Thomas & Young, 2010). Learning unfolds in many different ways within these environments, such as practicing intercultural negotiations through story-driven interactions with virtual characters (Traum, et al., 2008), or performing scientific inquiry within interactive mysteries where students play the roles of detectives (Ketelhut, et al., 2010).

### **1.1 Problem**

While narrative-centered learning environments show promise for enhancing student learning and engagement, students often have varying degrees of competency at solving ill-structured problems (Jonassen, 2000). When narrative-centered learning environments focus plots on complex problem-solving tasks—the types of problems that are critical for developing 21<sup>st</sup> century skills—students often struggle to advance the narratives unless effective problem-solving guidance is provided. Students approach games with significant individual differences in previous game-playing experience and content knowledge, which can have profound impacts on students' ability to learn and complete problem-centric story scenarios (Rowe, Shores, Mott, & Lester, 2010a). Consequently, when interacting with narrative-

centered learning environments, students often employ non-typical problem-solving methods or require context-specific instructional assistance to make substantive progress. A one-size-fits-all approach to interactive narrative and pedagogical design is ill suited to the requirements of narrative-centered learning. The motivational benefits of narrative-centered learning environments are likely to be lost if students flounder due to inadequate instructional support.

One-on-one, face-to-face human tutoring has long been considered the gold standard for effective instruction because of its significant pedagogical benefits compared to traditional classroom lectures. Work by the intelligent tutoring systems community has sought to emulate effective tutoring—such as personalized feedback, scaffolding, and problem sets—for a range of educational subjects (VanLehn, 2006; Woolf, 2008). However, most intelligent tutoring systems do not possess built-in mechanisms for reasoning about interactive narrative plots; the learning scenarios in intelligent tutoring systems typically abstract away narrative contexts in order to focus on educational subject matter. Embedding intelligent tutoring systems within narrative-centered learning environments requires integration with intelligent narrative technologies capable of dynamically generating and managing interactive stories (Mott & Lester, 2006b). However, research on integrating intelligent tutoring systems and narrative-centered learning environments is still in a nascent stage. To address this gap, this dissertation will examine the following research question: how can we devise effective technologies that dynamically tailor events in narrative-centered learning environments to support complex problem solving?

## **1.2 Approach**

In order to promote effective problem solving in narrative-centered learning environments, we investigate run-time systems that perform *narrative-centered tutorial planning*. Narrative-centered tutorial planning involves dynamically generating and revising story-centric problem-solving scenarios in accordance with educational goals and student needs. A wide

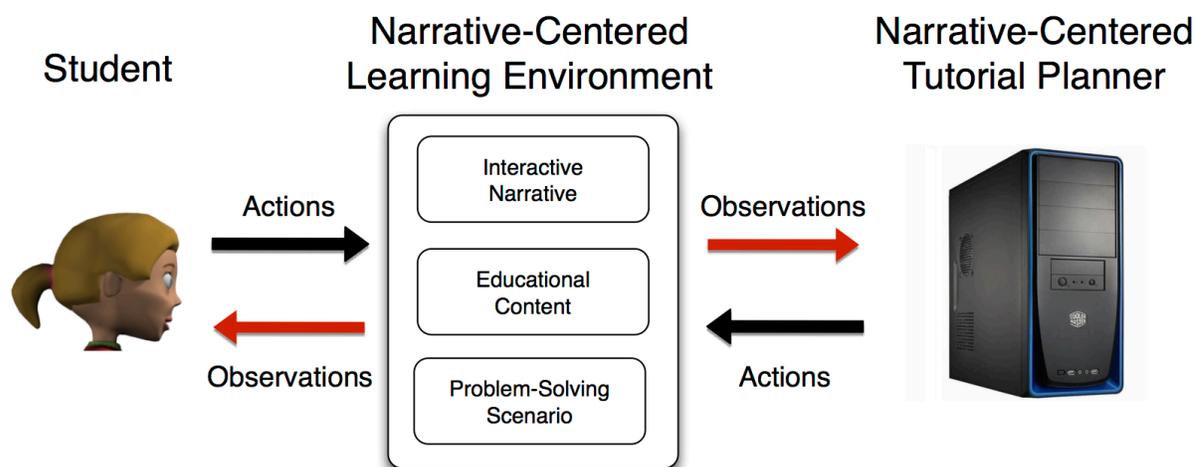


Figure 1. Student interaction with a narrative-centered tutorial planner.

range of story manipulations are possible in narrative-centered learning environments, such as dynamically augmenting plot sequences (introducing events, re-ordering events), modifying character states (revising character goals, adjusting character abilities), or presenting narrative-embedded instructional resources (delivering hints, introducing cognitive tools). Narrative-centered tutorial planners should be capable of performing a range of manipulations in order to achieve instructional objectives, such as scaffolding student problem solving, customizing scenario difficulty, or presenting alternate story experiences to enhance replayability and encourage time-on-task. A diagram depicting a student interaction with a narrative-centered tutorial planner is shown in Figure 1.

Narrative-centered tutorial planning unifies two related experience management problems: *pedagogical planning* and *drama management*. Pedagogical planners dynamically sequence learning exercises and deliver context-specific, tailored problem-solving supports in intelligent tutoring systems (VanLehn, 2006; Woolf, 2008). Drama managers perform dynamic adjustments to interactive narratives in order to heighten dramatic impact, foster narrative believability, or preserve authorial intent (Mateas & Stern, 2005; Riedl, Saretto, & Young, 2003; Roberts & Isbell, 2008). There has been some work leveraging drama

management technologies in narrative-centered learning environments (Aylett, et al., 2005; Si, Marsella, & Pynadath, 2005), but there are relatively few examples of systems that tightly integrate intelligent tutoring and computational models of interactive narrative (Thomas & Young, 2010; Mott & Lester, 2006b; Lee, Mott & Lester, 2012).

While narrative-centered tutorial planners show promise for addressing students' individualized needs, there is limited research available about how to design such systems effectively. As a consequence, data-driven approaches to devising narrative-centered tutorial planners are particularly attractive. Data-driven planners act in a principled manner—based on past observations of student behavior—while reducing the need for labor-intensive hand authoring of rules.

Narrative-centered tutorial planners should meet several operational requirements. First, they should perform instructional interventions through the affordances of stories. In other words, narrative-centered tutorial planners should primarily act by adjusting story events, character behaviors, and discourse elements, and these manipulations should take place within coherent and compelling plots. Second, narrative-centered tutorial planners should have a broad and diverse repertoire of story manipulations at their disposal. These capabilities are critical for maintaining flexibility in how planners aid students during various story contexts. Third, narrative-centered tutorial planners should explicitly address the multiple goals of narrative-centered learning, such as supporting complex problem solving, enhancing content learning, and fostering student engagement. Finally, narrative-centered tutorial planners must meet the real-time performance requirements of interactive narratives.

In order to address these requirements, this dissertation presents an empirical framework for inducing narrative-centered tutorial planners using concurrent Markov decision processes. Inspired by modular reinforcement learning techniques, the framework presents a principled method for decomposing sequential decision-making problems into multiple constituent parts (Bhat, Isbell, & Mateas, 2006; Sprague & Ballard, 2003). After decomposing the narrative-centered tutorial planning task, each component is formalized as an independent Markov decision process that is solved using certainty equivalent learning

techniques with previous students' data. The sub-problems are solved individually, and the solutions are combined to form an overall decision-making policy. The policies are explicitly trained to optimize designer-specified metrics that embody key goals of narrative-centered learning. The induced policies are ultimately re-incorporated into a narrative-centered learning environment to support new students' interactions with the system.

### **1.3 Thesis Statement and Hypotheses**

This dissertation investigates the following thesis statement:

*Modeling narrative-centered tutorial planning with concurrent Markov decision processes is an effective approach for dynamically tailoring story events to enhance student problem solving in narrative-centered learning environments.*

Assessing the quality and effectiveness of a narrative-centered tutorial planner is a multi-faceted problem. This investigation examines planner effectiveness by focusing on four principal requirements: 1) performing instructional interventions that are naturally embedded within an interactive narrative, 2) accommodating a broad range of potential story manipulations, 3) explicitly addressing the instructional goals of narrative-centered learning, and 4) satisfying the real-time performance requirements of interactive narratives.

Three instructional goals are considered in this work in order to characterize the quality of students' learning experiences. First, students' problem-solving performances are considered. Problem solving is assessed in terms of key educational sub-goals completed by students during narrative-centered learning experiences. Additionally, students' frequency and efficiency in completing overarching problem-solving tasks is measured. Second, students' content learning gains are considered. In this work, content learning gains are measured through curricular tests that are administered before and after instructional interventions. Third, students' engagement levels during and after narrative-centered learning experiences are considered. This work utilizes validated and reliable instruments for

measuring presence and intrinsic motivation to assess students' post hoc responses to narrative-centered learning environments.

In order to study the effectiveness of our framework for inducing narrative-centered tutorial planners with concurrent Markov decision processes, we investigate four primary hypotheses. We use the phrase *Induced Planner* to refer to a narrative-centered tutorial planner obtained through our framework.

- **Hypothesis 1:** Students who use a narrative-centered learning environment with an Induced Planner will solve complex problem-solving tasks more quickly and more frequently than students who use a baseline version of the environment.
- **Hypothesis 2:** Students who use a narrative-centered learning environment with an Induced Planner will demonstrate more deliberate problem-solving practices than students who use a baseline version of the environment.
- **Hypothesis 3:** Students who use a narrative-centered learning environment with an Induced Planner will achieve greater content learning gains than students who use a baseline version of the environment.
- **Hypothesis 4:** Students who use a narrative-centered learning environment with an Induced Planner will demonstrate greater engagement than students who use a baseline version of the environment.

## 1.4 Contributions

This dissertation reports on research that has made the following contributions:

- An implemented narrative-centered learning environment that reliably yields significant science content learning gains, as well as promising levels of student engagement (Rowe, Shores, Mott, & Lester, 2011; Sabourin, Rowe, Mott, & Lester, 2011).
- An empirical account of student learning outcomes, engagement outcomes, and gameplay characteristics, as well as their inter-relationships, in a narrative-centered learning environment (Rowe, McQuiggan, Robison, & Lester, 2009; Rowe, et al.,

- 2010a; Rowe, et al., 2011; Sabourin, et al., 2011; Spires, Rowe, Mott, & Lester, 2011).
- An empirical framework for devising narrative-centered tutorial planning policies that is based on concurrent Markov decision processes. This framework improves upon prior approaches to building narrative-centered tutorial planners by inducing narrative adaptation strategies directly from student interaction data.
  - This dissertation reports on the first *deployed* interactive narrative director agent whose narrative decision-making behavior is solely induced from user interaction data. The model does not require hand-authored rules about narrative decision-making strategies, or explicit demonstrations of narrative decision-making by humans.
  - A sizable corpus of student data from interactions with a narrative-centered tutorial planner. The corpus includes planner actions that are intentionally exploratory, and it can serve as a resource for a range of approaches to devising narrative-centered tutorial planners, as well as empirical studies of narrative-centered learning (Sabourin, et al., 2011; Sabourin, Rowe, Mott, & Lester, 2012).
  - A narrative-centered learning environment that dynamically adapts story events to enhance students' problem-solving processes. The planner is comprised of a suite of narrative adaptation policies, which are implemented as lookup tables mapping MDP states to planner actions, and meet the real-time performance requirements of interactive narratives. The system also supports easy modification of narrative adaptation policies for future investigations.
  - An empirical account of students' problem-solving processes, learning outcomes, and engagement outcomes when interacting with an adaptive narrative-centered learning environment imbued with an induced narrative-centered tutorial planner. Additionally, findings from an empirical evaluation of the induced narrative-centered tutorial planner are reported.

## **1.5 Organization**

The remainder of this dissertation is organized as follows. Chapter 2 provides background on narrative-centered tutorial planning. This includes a discussion of related work on narrative-centered learning environments, intelligent tutoring systems and interactive narrative technologies, as well as a detailed examination of decision-theoretic approaches to narrative adaptation. Chapter 3 presents our framework for inducing narrative-centered tutorial planners directly from student data using concurrent Markov decision processes. Chapter 4 describes CRYSTAL ISLAND, the test bed narrative-centered learning environment used to test our framework, as well as empirical results on student learning, engagement, and problem solving in a non-adaptive version of the environment. Chapter 5 describes an adaptive version of CRYSTAL ISLAND and an associated series of classroom studies used to collect a training corpus for automatically inducing the narrative-centered tutorial planner. Chapter 6 describes the implemented narrative-centered tutorial planner for CRYSTAL ISLAND, as well as an evaluation experiment conducted to examine its effectiveness in enhancing student problem solving. Chapter 7 presents results from the evaluation experiment, and Chapter 8 follows with a discussion of the findings' implications. Chapter 9 concludes the dissertation by revisiting our hypotheses and thesis statement, as well as discussing directions for future work.

## **CHAPTER 2**

### **Background and Related Work**

Recent years have seen a growing interest in integrating narratives and digital games for educational purposes. Research on narrative-centered learning environments is situated at the intersection of two fields: interactive narrative systems and intelligent tutoring systems. By combining techniques from intelligent tutoring systems and interactive narrative systems, narrative-centered learning environments can create adaptive story experiences that promote engaged learning and are tailored to individual students.

This chapter describes background and related work on dynamically tailoring story events in narrative-centered learning environments. Section 2.1 discusses the motivation for incorporating narratives in education, and it describes several existing narrative-centered learning environments. Section 2.2 describes work on intelligent tutoring systems, which use artificial intelligence (AI) technologies to dynamically coach and scaffold students' learning experiences. Section 2.3 surveys work by the intelligent narrative technologies community that is focused on dynamically adapting plots and discourses in interactive narrative systems. Section 2.4 describes several decision-theoretic frameworks for interactive narrative director agents, which share key features with the concurrent Markov decision process framework presented in this dissertation.

#### **2.1 Narrative-Centered Learning Environments**

Narratives are pervasive throughout human communication and culture. Narratives provide structure for encoding experiential knowledge and are an integral component in meaning making (Bruner, 1991; Polkinghorne, 1988). Graesser and Ottati (1996) argue that “stories have a privileged status in the cognitive system,” citing experimental findings that suggest readers process narrative texts more quickly and recall narrative information more readily

than expository forms. By extension, stories offer significant promise for enhancing learning and problem solving (Jonassen & Hernandez-Serrano, 2002). Stories are ubiquitous tools for sharing experiential knowledge, recounting prior problem solutions, and fostering vicarious experience. They can provide problem-solving guidance by serving as examples to be adapted to current situations. Additionally, stories are instrumental in assessment by virtue of their ability to present novel situations to test transfer of generalizable skills.

Despite the fact that narrative has long been believed to hold significant potential for education, it has traditionally been static rather than interactive, which has limited the types of narrative-centered learning experiences possible in educational settings. Only in recent years have advances in commercial game technologies enabled rich narratives to be integrated into virtual environments featuring cinematic presentations, believable characters, and dramatic plots. By repurposing these technologies for educational objectives, narrative-centered learning environments overcome the passive nature of traditional storytelling, while maintaining their motivational and pedagogical benefits.

There are three primary goals behind efforts to integrate narratives and digital games for education. The first goal is enhancing student motivation and engagement. Narratives and games are ubiquitous forms of entertainment, and work on narrative-centered learning environments aims to utilize their engaging qualities to evoke student interest and sustain learning. The second goal is to support students' meaning making processes, leveraging humans' natural facility for narratives to reinforce associations with educational material. The third goal is enhancing students' problem-solving skills. Given stories' role in supporting complex problem solving, narrative-centered learning environments are designed to guide students as they develop generalizable problem-solving skills.

Narrative-centered learning environments have been devised for a range of subject matters, including foreign language learning (Johnson, 2010), intercultural negotiation training (Kim et al., 2009; Traum, et al., 2008), anti-bullying education (Aylett, et al., 2005), computer security (Thomas & Young, 2010), and middle school science (Hickey, Ingram-Goble, & Jameson, 2009; Ketelhut, et al., 2010). Interactions with narrative-centered learning

environments can take several forms. Students may directly influence a narrative by completing actions in order to solve a problem, or they may indirectly influence events by providing guidance to autonomous virtual characters. Narrative-centered learning environments have been developed to support both single and multiple players, they have been realized using realistic 3D graphics engines as well as abstract cartoon-like representations, and they have structured problem-solving activities within overarching narratives, as well as sequences of related vignettes.

There have also been efforts to incorporate AI technologies into narrative-centered learning environments, including work on autonomous virtual characters (Aylett, et al., 2005; Kim et al., 2009; Marsella, Johnson, & Labore, 2003), natural language processing (Johnson, 2010; Kim et al., 2009), intelligent tutoring (Johnson, 2010; Kim et al., 2009; Thomas & Young, 2010), and narrative director agents (Marsella, Johnson, & Labore, 2003; Thomas & Young, 2010; Lee, Mott, & Lester, 2012). Following are descriptions of several narrative-centered learning environments.

**River City.** Developed at Harvard, River City is a multi-user virtual environment aimed at improving middle school students' deep inquiry skills and science content knowledge (Nelson, 2007; Ketelhut, et al., 2010). The game's narrative takes place in a late nineteenth century city whose residents have mysteriously fallen ill. Students control in-game avatars and work in teams to explore the virtual city, collect clues and evidence concerning the mysterious illness, formulate and test hypotheses, and compare research findings. Science content is integrated with historical, social, and geographical content. River City is accompanied by a 2–4 week in-class curriculum and is designed to replace traditional science lessons. Over the past decade, the software has been used in hundreds of classrooms with tens of thousands of students throughout the United States.

River City implements a rule-based guidance system to scaffold students' inquiry processes (Nelson, 2007). Individualized guidance messages are triggered by pre-selected events in the virtual environment, and they are tailored based on students' past actions in the

game. Several empirical studies of students' gameplay and learning data indicated that increased usage of the guidance system improved students' learning outcomes (Nelson, 2007). However, it was found that students with low self-efficacy for science used the guidance system less frequently than high self-efficacy students, and the guidance system was insufficient for closing the learning gap between low and high self-efficacy groups (Nelson & Ketelhut, 2008). The dynamic guidance system in River City focuses solely on tutorial support as opposed to narrative experience, and therefore comprises a restricted class of narrative adaptations: non-diegetic cognitive prompts. In other words, tutorial guidance is not naturally embedded within the interactive narrative. Instead, tutorial messages appear to come from a source external to the storyworld. Further, the delivery of guidance messages is based on extensive human-authored rules.

**Quest Atlantis.** Quest Atlantis is a narrative-centered, multi-user virtual environment developed at Indiana University that has been used by over fifty thousand students internationally (Barab, Gresalfi, & Ingram-Goble, 2010; Hickey, Ingram-Goble, & Jameson, 2009). The software features a complex storyline about the fictional world of Atlantis. The Atlantians' planet is in rapid decline, and students must help to restore lost Atlantian knowledge that has precipitated the world's social and environmental decay. Gameplay activities are distributed across several virtual worlds. The virtual worlds feature distinct problem-solving scenarios that connect to national and local academic standards. For example, the Taiga Park world focuses on a riverside community with a declining fish population (Hickey, Ingram-Goble, & Jameson, 2009). In Taiga Park, students complete a series of quests that incorporate socioscientific inquiry and ecological science concepts, incrementally addressing the community's looming ecological and economic dilemma. Students interact with virtual characters, collect and analyze data, and write and submit reports in order to remediate the river.

Quest Atlantis's learning benefits have been investigated for several academic subjects, including middle school science (Hickey, Ingram-Goble, & Jameson, 2009) and

language arts (Barab et al., 2010; Warren, Dondlinger, & Barab, 2008). Empirical studies with a range of student populations have found that Quest Atlantis provides significant learning and engagement benefits compared to traditional, text-based classroom activities. However, Quest Atlantis's core research team has not investigated AI technologies for dynamically tailoring students' learning experiences. Recent work by Shute and colleagues (2010) has proposed a Bayesian network framework for stealth assessment in narrative-centered learning scenarios like Taiga Park, but to our knowledge it has not been implemented and evaluated in the Quest Atlantis software.

**Carmen's Bright IDEAS.** Developed at the University of Southern California, Carmen's Bright IDEAS is an interactive pedagogical drama designed to teach social problem-solving skills to mothers of pediatric cancer patients (Marsella, Johnson, & Labore, 2003; Marsella, Johnson, & LaBore, 2000). Learners influence the choices of the main character, Carmen, as she copes with the difficulties inherent in caring for a chronically ill child. Carmen is given coaching by Gina—a second autonomous agent—to implement the Bright IDEAS approach to social problem solving. The environment's autonomous virtual characters implement multi-layer state-machine dialogue models, models for gesture and non-verbal behavior, and an appraisal-based model of emotion, to foster dramatic interactions between the narrative's two characters. By carefully coordinating the dialog, emotional, and non-verbal behaviors of the Gina and Carmen animated agents, as well as decisions made by director and cinematographic agents, Carmen's Bright IDEAS aims to balance the demands of narrative, interactivity, and learning. A prototype version of the system was investigated through exploratory clinical trials involving twenty-six human participants (Marsella, Johnson, & Labore, 2003). The participants responded positively, reporting that the narrative seemed believable and was helpful for understanding how to apply the Bright IDEAS framework.

While Carmen's Bright IDEAS provides dynamic problem-solving support through character behaviors and non-diegetic text prompts, their delivery is based on hand-authored behavior rules. Authoring these behaviors, which should simultaneously promote compelling

drama and effective learning in a range of contexts, introduces considerable challenges. Furthermore, the system's narrative adaptations are primarily focused on believable and dramatic character performances; they are not explicitly designed to maximize student learning gains and engagement. To our knowledge, no empirical findings have been reported about the learning and engagement effects of the system's dynamic narrative adaptations.

**FearNot!** FearNot! is a character-driven, interactive narrative learning environment that promotes anti-bullying social education (Aylett, et al., 2005; Vannini, et al., 2011). The system uses affectively-driven autonomous agents to dynamically generate dramatic, educational vignettes about bullying. In between non-interactive vignettes, virtual bullying victims consult the student for advice on coping strategies, and then use this feedback to inform their behavior in subsequent vignettes. The system's autonomous agents use an emotion model that implements a double appraisal process to predict the affective impact of actions on other agents (Aylett & Louchart, 2008). The double appraisal process serves as a metacognitive reasoning step for guiding agent behaviors to produce interesting narrative experiences. FearNot! uses an emergent narrative approach to story generation; the technique is contrasted with top-down interactive narrative generators that leverage explicit plot models to drive story generation and adaptation.

FearNot! has been the subject of large-scale, cross-cultural evaluations, but it has not been observed to significantly impact students' knowledge about bullying or coping strategies (Watson et al., 2010). While the system's emergent approach to narrative generation is likely to produce believable and dramatic narrative experiences, its use of minimally interactive bullying vignettes reduces user control and agency. Further, the system provides little explicit tutorial support to enhance learning gains.

**Tactical Language and Culture Training System.** The Tactical Language and Culture Training System (TLCTS) is a suite of story-centric, serious games designed for language and culture learning, and it is currently in use by tens of thousands of learners (Johnson,

2010). Primarily developed at Alelo Inc., TLCTS uses a combination of interactive lessons and narrative scenarios to train culturally embedded spoken and non-verbal communication skills. The Skill Builder component of the system provides interactive instruction about relevant cultural and language content. The Mission Game component provides opportunities to apply communication skills through task-based narrative scenarios involving face-to-face interactions with virtual characters. TLCTS has been the subject of several in vivo studies involving hundreds of military personnel (Johnson, 2010). These evaluations have found that TLCTS consistently elicits significant language and culture learning gains, as well as self-efficacy and motivational benefits.

TLCTS uses AI technologies extensively throughout its more than one hundred hours of training content; AI serves a key role in lesson authoring, automated speech recognition, virtual character dialogue and non-verbal behavior, tutorial feedback, and student modeling facilities. While initial versions of TLCTS also incorporated AI technologies for interactive narrative adaptation (Si, Marsella, & Pynadath, 2005), recent versions have forgone this functionality.

**BiLAT.** BiLAT is a serious game developed at the USC Institute for Creative Technologies that enables students to practice cross-cultural negotiation skills (Kim et al., 2009). Learning activities in BiLAT revolve around a series of bilateral negotiations with virtual Iraqi citizens, such as a hospital administrator, police chief, and doctor. Students prepare for each negotiation by completing a leader preparation worksheet, in which the student organizes background information and beliefs that are relevant to the forthcoming meeting. The actual negotiations are realized through multi-modal conversations with the virtual characters. Students advance the negotiations by performing sequences of menu selections, with action categories such as ‘say,’ ‘ask,’ ‘give,’ and ‘do.’ Implicit feedback about students’ negotiation decisions is delivered through character responses during the meeting; explicit feedback is provided in the form of coaching messages and guided after-action reviews. Initial studies of the BiLAT system found learning benefits for users with little prior negotiation experience

(Kim et al., 2009). However, no discernible benefits were observed for users with substantial prior negotiation experience. Challenges associated with assessing ill-defined knowledge (such as negotiation skills) point to the need for further research.

BiLAT leverages several AI technologies for dynamically tailoring students' learning experiences. The system implements a rule-based dialogue model for driving virtual characters' conversational dynamics, as well as a negotiation model based on the PsychSim simulation tool (Pynadath & Marsella, 2005). Intelligent tutoring facilities deliver individualized guidance in the form of hints and feedback, as well as structured after-action reviews. Narrative adaptations are achieved through variations in virtual character behaviors and non-diegetic coaching messages. To our knowledge, empirical evaluations have not yet been conducted that isolate the individual impacts of the system's intelligent tutoring or character behavior components.

**FixIt.** FixIt is an exploratory narrative-centered learning environment about computer security (Thomas, 2011; Thomas & Young, 2010). Students navigate a 3D virtual environment that represents a fictional computer. The virtual environment includes elements that denote the computer's CPU, processes, buses, and data packets. Students are responsible for identifying and repairing rogue processes that endanger the computer's operation. FixIt was implemented as a test bed environment for investigating Annie, a narrative-centered tutorial planner that delivers customized scaffolding for task-oriented, procedural domains. Annie utilizes a decompositional partial order planning framework for knowledge representation and reasoning, as well as diagnostic and remediation templates for pedagogical scaffolding. Annie extends existing narrative planning techniques—previously applied in entertainment-centered interactive narratives (Riedl, Saretto, & Young, 2003)—in order to incorporate tutorial strategies in interactive narrative decision making. The system maintains a course-grained student model, and it uses the model to drive story elaborations that are intended to achieve student knowledge objectives.

Annie's tutorial and narrative adaptations provide guidance about the preconditions and effects of important problem-solving actions in a learning environment (Thomas, 2011). The narrative adaptations performed by Annie provide leveled coaching that observes the "prompt, hint, show, teach, do" stages of pedagogical assistance. Annie's remediation actions can be realized in three forms: text-based messages, short cinematic sequences, and concrete actions that are performed by temporarily taking control of the player character. A pilot study found that Annie effectively predicts students' self-assessments of knowledge after they interact with the FixIt narrative-centered learning environment (Thomas & Young, 2010). A second study found that Annie also enhances student task completion in the FixIt environment (Thomas, 2011). However, the two studies did not find evidence that Annie enhances students' learning and engagement outcomes. This may have been due to the short duration of the learning experiences, and limited statistical power stemming from small numbers of study participants.

Several of the narrative-centered learning environments described in this section employ AI technologies to dynamically manage students' learning experiences. However, most of these systems loosely couple tutorial scaffolding with narrative adaptation decisions; explicit pedagogical supports are rarely delivered through naturalistic story structures. Furthermore, all of these systems sequence pedagogical supports based on hand-authored knowledge structures; none use empirical models to drive narrative and tutorial decision making. The nascent field of narrative-centered tutorial planning research has not yet developed a set of best practices for narrative-centered learning environment design. Given the dearth of theoretical and practical guidance, employing data-driven approaches to narrative decision making is a promising approach.

Devising effective pedagogical models using machine learning techniques has been an active area of research within the intelligent tutoring systems community for several years. The next section surveys work on intelligent tutoring systems, with emphases on common system behaviors and computational representations.

## 2.2 Intelligent Tutoring Systems

An intelligent tutoring system is a type of educational software that uses explicit models of students' learning processes in order to dynamically tailor instruction and pedagogical assistance to individual learners. Using one-on-one human tutoring as a model, intelligent tutoring systems monitor students' cognitive, affective, and metacognitive states in order to revise instructional plans, manage educational scaffolding, and deliver empathetic support (Shute & Psotka, 1996; Woolf, 2008). Intelligent tutoring systems have been developed for a broad range of academic subjects, including mathematics (Koedinger, Anderson, Hadley, & Mark, 1997), science (Graesser et al., 2004; VanLehn et al., 2005), writing (Britt, Wiemer-Hastings, Larson, & Perfetti, 2004), and language learning (Johnson, 2010), as well as training tasks, such as aircraft troubleshooting (Mislevy & Gitomer, 1995) and naval ship systems operation (Rickel & Johnson, 1999). Intelligent tutors have also been shown to provide significant educational benefits, yielding learning gains that are up to one standard deviation unit better than traditional classroom instruction (Koedinger, et al., 1997; VanLehn et al., 2005).

Over the past decade, general frameworks have emerged for characterizing the behavior of intelligent tutoring systems (VanLehn, 2006). According to VanLehn (2006), intelligent tutors operate in terms of two loops: an outer loop and an inner loop. The outer loop is responsible for making coarse-grained adjustments to the direction of a tutoring experience. The loop iterates over learning tasks, and strategically sequences instructional opportunities based on students' abilities. For example, the outer loop de-emphasizes skills that a student has already mastered, and increases time on topics requiring additional practice. The inner loop is responsible for delivering tailored feedback and hints during problem solving. Inner loop adaptations typically address a single educational exercise, feedback is context-sensitive, and pedagogical support is tailored to individual students' needs.

A variety of techniques have been employed for knowledge representation and reasoning in intelligent tutors. Symbolic representations such as production rules, semantic

networks, and constraint pairs have been especially prevalent. The Cognitive Tutor family of systems represents students' procedural knowledge using cognitive science-inspired knowledge bases of if-then rules (Anderson, Boyle, Corbett, & Lewis, 1990). The knowledge models are based on the ACT-R cognitive architecture, and they store comprehensive, fine-grained representations of students' domain knowledge to drive context-specific feedback and hints. The Stat Lady intelligent tutor leverages hierarchical semantic networks of student knowledge, which specify instructional relationships between symbolic, procedural, and conceptual knowledge components to guide content sequencing (Shute, 1995). SQL Tutor uses a constraint-based approach to knowledge representation that avoids direct storage of comprehensive knowledge models (Mitrovic & Ohlsson, 1999). Instead, SQL Tutor focuses on characteristics of completed problem solutions, defined as constraint pairs, in order to facilitate error identification and remediation during problem solving.

Symbolic approaches remain common in intelligent tutoring, but statistical representations have become increasingly prevalent for pedagogical reasoning and assessment. Probabilistic graphical models have been incorporated into several intelligent tutors in the forms of Bayesian belief networks. One example is Andes, which automatically generates problem-specific Bayesian networks to drive plan recognition and knowledge tracing in a tutor for college-level physics (VanLehn et al., 2005). Bayesian networks provide a principled representation for reasoning about the uncertainty inherent in models of students' plans and knowledge. Andes' probabilistic assessments are used to select tailored hints and feedback that are appropriate for students' inferred problem-solving plans. In a successor project called DT Tutor, Andes' Bayesian networks were extended into dynamic decision networks, tightly integrating student modeling and pedagogical decision making through a unified decision-theoretic framework (Murray, VanLehn, & Mostow, 2004).

The probabilistic representations used by Andes and DT Tutor are compatible with machine learning algorithms. However, many of the models' parameters were hand-authored. In contrast, intelligent tutors have also leveraged machine learning techniques to induce models from student data. Hidden Markov models have seen broad adoption by the

intelligent tutoring systems community for a range of tasks. For example, hidden Markov models are used for knowledge tracing throughout the Cognitive Tutor suite of educational systems (Corbett & Anderson, 1994). The Java Tutor system leverages hidden Markov models to emulate conversational patterns during tutorial dialogues (Boyer et al., 2010). The models are trained from corpora of human-human tutoring sessions in which students work with a tutor to solve introductory programming exercises. Jeong et al. (2008) use hidden Markov models to identify students' learning behavior patterns in the Betty's Brain learning-by-teaching environment.

As an alternative to probabilistic graphical models, reinforcement learning techniques have been the subject of growing interest in the intelligent tutoring systems community (Barnes & Stamper, 2008; Chi, VanLehn, & Litman, 2010). This work has emphasized probabilistic models of behavior, as opposed to explicit models of cognitive states, in order to analyze student learning. In reinforcement learning approaches, pedagogical decisions about tutorial feedback and hint selection are formalized as Markov decision processes. Tutorial policies are induced in order to optimize a delayed reward signal, which is usually based on students' learning outcomes. ADVISOR is one of the earliest examples of an intelligent tutoring system based on reinforcement learning (Beck, Woolf, & Beal, 2000). The system used temporal-difference learning to induce teaching policies from corpora of simulated student data. ADVISOR was able to devise teaching policies that reduced students' problem-solving times and generalized across distinct student populations for a grade school mathematics tutor. More recently, Chi, VanLehn and Litman (2010) have used Markov decision processes to model students' tutorial dialogues, devising pedagogical tutorial tactics directly from student data in the Cordillera physics tutor. Barnes and Stamper (2008) modeled students' logic proof sequences as Markov decision processes in order to automatically generate context-appropriate hints. Barnes's and Stamper's approach substantially reduces the authorial burden associated with hand crafting hint models by employing historical student data.

Intelligent tutoring can be viewed as an instance of the experience management problem (Riedl, Stern, Dini, & Alderman, 2008). An alternative experience management task is dynamically adapting story structures in interactive narrative systems. While the two tasks raise distinct challenges and objectives, they also share a common emphasis on dynamically augmenting software interactions in order to enhance users' experiences. The intersection of these two areas is central to the work presented in this proposal. The next section provides a high-level summary of computational issues and approaches for dynamically augmenting interactive narrative experiences.

### **2.3 Dynamic Adaptation of Interactive Narratives**

Interactive narrative systems are a type of digital storytelling media in which users play active roles in ongoing narratives. A key objective of interactive narrative systems is promoting user agency (Murray, 1998). This emphasis on agency is predicated on the assumption that accommodating impactful actions in virtual story worlds enhances users' sense of narrative transportation.

In order to promote user agency, the intelligent narrative technologies community has investigated AI techniques for dynamically adapting interactive stories in response to users' actions. One important line of investigation in interactive narrative systems is balancing user agency with story coherence and authorial intent. As described by Riedl, Saretto, and Young (2003), when users perform actions that have meaningful impacts on interactive stories, their actions may also disrupt the "intended" narrative experiences desired by authors or other virtual agents. These conflicts can result in coherence and cohesion breakdowns. Intelligent narrative technologies can be utilized to preserve story coherence by intervening or responding to unpredictable user actions (Riedl, Saretto, & Young, 2003). An additional focus of interactive narrative research is varying story experiences in order to encourage novelty and re-playability (Roberts, et al., 2006). Rather than focus on delivering high-quality but homogeneous story experiences, these systems aim to produce targeted distributions of story experiences over multiple sessions. Furthermore, work on user-adaptive

digital games has investigated dynamically altering the presentation and content of story events so that they are tailored to individual users' preferences (Thue, Bulitko, Spetch, & Wasylishen, 2007).

Architectures for performing interactive narrative planning are often characterized in terms of story-centric approaches and character-centric approaches (Porteous, Cavazza, & Charles, 2010; Si, Marsella, & Riedl, 2008). Story-centric approaches use explicit models of stories in order to guide interactive narrative generation. Decision-making capabilities are localized in an omnipotent director agent (or drama manager) that maintains a global perspective of the story and reasons about the state and progression of the plot (Mateas & Stern, 2005; Riedl, Saretto, & Young, 2003; Roberts, et al., 2006). Alternatively, narrative planning can be decentralized by employing character-based approaches to interactive narrative generation. This family of approaches emphasizes autonomous agents that act in a natural and believable manner within a story environment (Aylett, et al., 2005; Cavazza, Charles, & Mead, 2002; Si, Marsella, & Pynadath, 2005). In this approach, narratives emerge through interactions between agents and the user. This approach is well suited for producing believable character dynamics, but it often lacks mechanisms for ensuring particular plot sequences occur or preserving authorial intent. Hybrid approaches have also been investigated, which combine centralized director agents and virtual agents that perform locally autonomous behaviors (Riedl, et al., 2008). The narrative-centered tutorial planning framework presented in this dissertation adopts an alternative hybrid: a decentralized story-centric approach. The planner maintains an omnipotent view of the interactive narrative, but the narrative-centered tutorial planning task is decomposed into independent sub-problems.

Several families of algorithms have been employed in story-centric interactive narrative systems. Classical planning techniques are one prevalent approach, because STRIPS-style plans align naturally with computational representations of plots and discourses (Young, 1999). Traditional STRIPS planners produce operator sequences that achieve well-defined goal states (Fikes & Nilsson, 1971). In interactive narrative systems, this corresponds to sequences of plot events that lead to particular story conclusions. Plan-

based director agents monitor and revise the executions of story plans in order to respond to users' actions and preserve desirable narrative properties such as story coherence (Riedl, Saretto, & Young, 2003). In addition to classical planning techniques, reactive planners have been investigated for dynamically responding to user actions under real-time performance constraints. Several of these systems incorporate special-purpose data structures inspired by narrative concepts, such as dilemmas or beats, in order to bundle story content for reactive delivery (Barber & Kudenko, 2007; Mateas & Stern, 2005).

Search-based approaches have been investigated for dynamically managing interactive narratives. Search-based approaches attempt to find plot sequences that optimize designer-specified evaluation functions (Nelson & Mateas, 2005; Weyhrauch, 1997). These formalisms often use memoization or depth-bounded search techniques in order to constrain their computation times. However, they are sensitive to the narrative evaluation functions employed, which may be difficult to craft for a range of artistic aesthetics.

Case-based reasoning techniques have been used in several story-centric interactive narrative systems. The OPIATE story director dynamically responds to user actions by retrieving Proppian sub-plots rooted in particular story contexts (Fairclough, 2004). Sub-plot suitability is determined by using a k-nearest neighbor technique that considers character attitudes and storyworld states. In situations where a suitable case cannot be retrieved, multiple cases are combined to form a new sub-plot for delivery in the interactive narrative. Work by Sharma et al. (2007) modifies the search-based drama management approach of (Weyhrauch, 1997) by incorporating a case-based player model that approximates users' plot preferences. The cases are based on prior users' interactions with the narrative environment, and a drama manager uses the preference model to inform its narrative adaptation decisions.

More recently, statistical natural language processing techniques have been used to generate text-based interactive narratives. SayAnything is an interactive narrative system that leverages off-the-shelf information retrieval techniques to retrieve and present narrative fragments from a large corpus of online blogs (Swanson & Gordon, 2008). Users collaboratively author a text-based story with the system. After the user provides a sentence,

SayAnything identifies a corresponding sentence from the corpus and presents it as a new addition to the emerging story. SayAnything differs from many interactive narrative systems in that it does not seek to achieve particular authorial objectives. Instead, the system dynamically creates original story experiences. Additionally, its adoption of corpus-based approaches eliminates the authorial burden and domain-specificity constraints of many interactive narrative systems.

Recently, Lee, Mott, & Lester (2012) introduced a supervised machine learning framework for controlling interactive narrative director agents by training models from human demonstrations. This work involves conducting “Wizard of Oz” studies, where human users control embodied director agents and interact with other users in the storyworld, to collect corpora for training director agents. The human wizards’ directorial decisions become labels for machine learning strategies for how and when to perform narrative adaptations. An investigation of Dynamic Bayesian networks for modeling interactive narrative director agents found that the machine-learned model was highly effective at enhancing users’ interactive narrative experiences compared to baseline director agent techniques. However, the logistical challenges inherent in conducting Wizard of Oz studies are a notable downside of this approach; wizards require advance preparation, and they can become a major bottleneck while collecting training data. While more real-world variations of this design are possible (e.g., players are randomly paired during interactions in an online environment, one serving as the protagonist and the other serving as the director agent), such variations have not yet been systematically investigated.

Another important class of narrative adaptation techniques relies on decision-theoretic planning algorithms. This category is most closely related to the concurrent Markov decision process framework that is the focus of this dissertation. A more detailed summary of this work is presented in the next section.

## **2.4 Decision-Theoretic Approaches for Narrative Adaptation**

Some interactive narrative systems have begun to use decision-theoretic techniques for

selecting narrative adaptations (Mott & Lester, 2006a; Nelson, Roberts, Isbell, & Mateas, 2006; Si, Marsella, & Pynadath, 2009). Several features typify decision-theoretic approaches to interactive narrative planning. First, they leverage explicit models of narrative utility and reward to drive reasoning. Second, they forecast future story events in order to compare the consequences of different courses of action. Third, they model the inherent uncertainty in user interactions with virtual story worlds. By reasoning about narrative events' relative values, as well as probabilities associated with those estimates, decision-theoretic approaches provide principled methods for selecting among competing narrative planner actions.

Evaluation is an essential aspect of decision-theoretic interactive narrative systems. Assessments often rely on objective evaluation functions or narrative-based metrics. Viewing narrative through an “objective lens” can be controversial because subjectivity and aesthetics are integral to artistic story experiences. However, decision-theoretic methods provide a useful paradigm for rational decision making, and they are particularly useful for interactive story environments with extrinsic objectives such as education.

U-DIRECTOR is an example of a decision-theoretic director agent that was implemented in an early version of the CRYSTAL ISLAND narrative-centered learning environment (Mott & Lester, 2006a). The director agent triggers story-based hints that assist students while they investigate an interactive science mystery. U-DIRECTOR reasons about multiple aspects of an interactive narrative experience—plot state, pedagogical state, and student state—using dynamic decision networks for recurrently selecting and evaluating directorial actions. During each time step, the director observes the current state of the narrative-centered learning experience, updates its decision network, and selects an action that maximizes expected narrative utility. While U-DIRECTOR's dynamic decision networks are theoretically compatible with machine learning approaches, in practice the model is automatically generated by hand-authored rules due to practical constraints associated with learning complex decision network structures and parameters.

THESPIAN is another example of a decision-theoretic interactive narrative system, which was implemented in a pedagogical drama for language learning (Si et al., 2009; Si,

Marsella, & Pynadath, 2005). THESPIAN endows virtual characters with goal-oriented decision-making models that are loosely based on partially observable Markov decision processes. Characters' goals are defined numerically and denote degrees of achievement for particular narrative states. This arrangement provides a means for defining characters' narrative and pedagogical priorities, and by extension their personalities, in order to guide narrative decision making. Additionally, each virtual character implements a recursive "theory of mind" model in order to reason about how its actions impact the beliefs and goals of other agents. One particularly attractive characteristic of THESPIAN is that characters' goals can be automatically configured by "fitting" them to a corpus of linear scripts. This reduces the authorial burden associated with crafting virtual characters that fill particular roles within a story or educational experience. However, unlike other machine learning approaches for devising interactive narrative models, the automated fitting process does not explicitly optimize users' narrative or educational experiences. Therefore, the system provides few guarantees about the quality of the generated interactive narratives.

More than other computational frameworks, declarative optimization-based drama management (DODM) is closely related to this dissertation's approach to narrative-centered tutorial planning. Primarily investigated by researchers at the Georgia Institute of Technology (Nelson, et al., 2006; Roberts, et al., 2006), DODM encompasses a group of interactive narrative techniques that model drama management as a Markov decision process. DODM encompasses models based on classical reinforcement learning (Nelson, et al., 2006) as well as targeted trajectory distribution Markov decision processes (Roberts, et al., 2006). Key components of the DODM framework include: discrete state representations that are defined in terms of plot point sequences, finite drama manager action sets, and hand-crafted reward functions that assess the aesthetic qualities of narratives or desired probability distributions over narrative experiences. DODM models' parameters are automatically induced using on-line reinforcement learning techniques (such as temporal-difference learning) with large interactive narrative corpora generated from simulated users.

The framework investigated in this dissertation is inspired by work on modular

reinforcement learning, and it shares several representational characteristics with DODM. Both frameworks model interactive narrative adaptation in terms of Markov decision processes. Both frameworks leverage data-driven techniques to induce model parameters, and they both use explicit objective functions to guide this process. However, the proposed narrative-centered tutorial planning framework differs from DODM in several important ways. Existing DODM implementations centralize directorial control within a single monolithic drama manager. The current proposal presents a framework that decentralizes narrative-centered tutorial planning in a plot-centric manner. Decentralization is achieved by decomposing the decision-making task into independent sub-problems. Solutions are obtained for each sub-problem, and they are combined through arbitration procedures in order to comprise a unified narrative-centered tutorial planner.

An additional distinction between the two frameworks is the source of training data they use. Existing DODM implementations are data intensive, and they have traditionally relied on corpora generated by simulated users. In contrast, the framework presented in this dissertation emphasizes modularity and constrained Markov decision process definitions, which enable usage of smaller corpora. These constraints make it possible to induce narrative-centered tutorial planning policies directly from student data, thereby reducing (or eliminating) the need for simulated users. This is particularly attractive for complex domains with ill-understood dynamics and multiple simultaneous objectives such as narrative-centered learning.

In this chapter, we have provided background on two fields that lay the foundation for this dissertation: interactive narrative technologies and intelligent tutoring systems. We have examined narrative-centered learning environments for a range of educational subjects, and analyzed their capacities to adaptively support students' problem solving and learning processes. Furthermore, we have discussed related work on AI techniques for dynamically tailoring story events in interactive narratives. Among these techniques, we have paid particular attention to decision-theoretic approaches for narrative adaptation, because they share several notable properties with the framework presented in this dissertation. Now that

we have established the necessary background, the next chapter presents our framework for dynamically tailoring story events in narrative-centered learning environments with concurrent Markov decision processes.

## CHAPTER 3

### Inducing Narrative-Centered Tutorial Planners

This chapter introduces a general framework for formalizing narrative-centered tutorial planning in terms of concurrent Markov decision processes, as well as techniques for addressing key issues that arise from this modeling process. Section 3.1 provides background on modular reinforcement learning, which provides the conceptual basis for the framework. Section 3.2 discusses the motivation for representing narrative-centered tutorial planning in terms of concurrent Markov decision processes, as well as operational details about this formalization. Section 3.3 discusses data-driven techniques for inducing narrative-centered tutorial planners in this framework. This includes a discussion of approaches for inducing effective narrative-centered tutorial planning policies, as well as employing arbitration techniques for combining modularized policies. Section 3.4 provides a theoretical analysis of the framework, examining both training corpus size requirements, as well as time complexity during model induction.

#### 3.1 Modular Reinforcement Learning

Modular reinforcement learning is a multi-goal extension of classical single-agent reinforcement learning (Bhat, Isbell, & Mateas, 2006; Karlsson, 1997; Sprague & Ballard, 2003). *Reinforcement learning* refers to a class of machine learning problems that involve delayed rewards (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). In classical reinforcement learning, an agent must learn a policy for selecting actions in an uncertain environment in order to accomplish a goal (or several goals). The environment is characterized by a set of states and a probabilistic model governing transitions between those

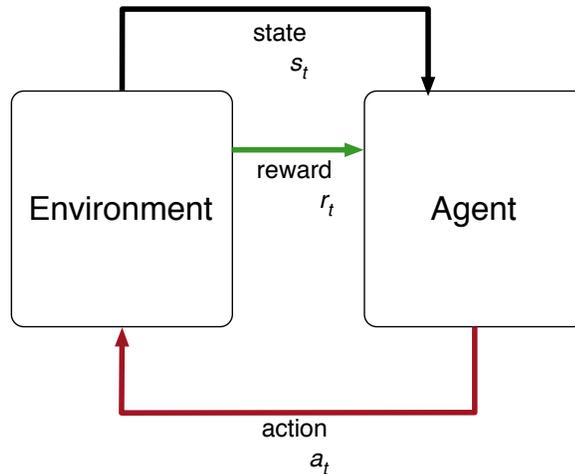


Figure 2. Agent-environment architecture for classical reinforcement learning.  
Adapted from (Sutton & Barto, 1998).

states. The agent is capable of observing the environment's state and using its observations to guide decisions about which actions to perform. In contrast to supervised machine learning, the agent is not provided with external instruction about which actions to take. Instead, the environment produces rewards that provide positive or negative feedback about the agent's actions. This feedback indicates whether the actions contribute toward achievement of the desired goal. The agent's task is to utilize the reward signal in order to learn a policy, denoted  $\pi$ , which maps observed states to actions and maximizes its total accumulated reward. Figure 2 illustrates the relationships between an agent, an environment, actions, states, and rewards in a classical reinforcement learning task.

A reinforcement learning problem is typically formalized as a *Markov decision process* (MDP) (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). Solution methods for MDPs typically assume that the decision-making task satisfies the Markov property, which requires that future events in the environment be independent of past events, given the environment's current state and the agent's current action. An MDP is formally defined by a tuple  $M = (S, A, P, R)$ , where  $S$  is the set of environment *states*;  $A$  represents the

set of *actions* that the agent can perform;  $P$  is the *state transition model*,  $P: \{S \times A \times S\} \rightarrow [0, 1]$ , which specifies the probability of transitioning to state  $s_{t+1} \in S$  after performing an action  $a_t \in A(s_t)$  in state  $s_t \in S$  at time step  $t$ ; and  $R$  is the *reward model*,  $R: \{S \times A \times S\} \rightarrow \mathcal{R}$ , which specifies the expected scalar reward  $r_t \in \mathcal{R}$  associated with performing action  $a_t$  in state  $s_t$  and transitioning to state  $s_{t+1}$ . The solution to an MDP is an optimal policy,  $\pi^*(s_t) \rightarrow A$ , that maps states to actions and yields maximum expected reward for the agent<sup>1</sup>.

In the classical version of reinforcement learning, the environment is *fully observable*, states and actions are both *finite* and *discrete*, the environment is *static* (i.e. the probabilities in the state transition model do not change), and proximal rewards possess greater value than distal rewards of equivalent magnitude. Additionally, reinforcement learning problems can be *episodic*, where an agent’s action sequences are divided into multiple independent episodes, or *continuing*, where the agent performs actions without regard for a terminal condition. It should be noted that the formulation described here represents the most elementary form of a reinforcement learning task. Considerable work has investigated problems that involve relaxing the aforementioned assumptions (Doya, 2000; Jaakkola, Singh, & Jordan, 1995; Littman, 1994), but these extensions are not essential for understanding our framework for narrative-centered tutorial planning.

Over the past decade, reinforcement learning techniques have been investigated for a range of complex applications, such as robotic soccer (Riedmiller, Gabel, Hafner, & Lange, 2009; Stone, Sutton, & Kuhlmann, 2005), automated vehicle control (Abbeel, Coates, Quigley, & Ng, 2007), and game playing (Amato & Shani, 2010; Sharma, Holmes, et al., 2007). Many real-world reinforcement learning problems require large training data sets. In these cases, simulations are the primary approach for gathering adequate training data.

---

<sup>1</sup> Notation for the formal description of MDPs is borrowed from (Sutton & Barto, 1998) and (Roberts, et al., 2006).

Reinforcement learning frameworks have also been investigated in intelligent user interfaces, such as dialogue systems (Chi, VanLehn, & Litman, 2010; Singh, Litman, Kearns, & Walker, 2002) and intelligent tutors (Barnes & Stamper, 2008; Beck et al., 2000). Human users play a critical role in determining the dynamics of these systems, but the inherent uncertainty associated with modeling human users introduces notable challenges in comparison to physical systems. Several efforts to utilize reinforcement learning for intelligent user interfaces have eschewed simulation-generated corpora in favor of corpora comprised of human user data (Chi, VanLehn, & Litman, 2010; Singh, Litman, Kearns, & Walker, 2002). However, collecting human user data is costly. These investigations have required steps to abstract reinforcement learning tasks in terms of approximate representations, trading optimality in exchange for learning efficiency.

Ideally, reinforcement learning models use state representations that encompass the agent's full history and include all state features that are relevant for decision making. However, reinforcement learning methods suffer from the *curse of dimensionality*, which holds that the problem's state space grows exponentially in the number of state features considered. One approach to addressing the curse of dimensionality is by mapping true MDP states to *approximate states*, which are often comprised of small sets of carefully selected state features (Singh, Kearns, Litman, & Walker, 1999). These features are often coarsely discretized. In some cases, approximate states may not even be strictly Markovian, but they often yield good policies in practice by approximating the Markov property.

A second, complementary approach for applying reinforcement learning to intelligent user interfaces is modular reinforcement learning, which factors complex state spaces by decomposing high-level tasks into multiple independent, concurrent sub-problems (Barto & Mahadevan, 2003; Karlsson, 1997; Sprague & Ballard, 2003). Models for each sub-problem are individually learned, and they need only consider those state features, actions and goals that are relevant to the sub-problem. The learned policies can be used in parallel during runtime, and arbitration techniques are employed when individual policies recommend competing actions. Solutions to modular reinforcement learning problems are typically

approximations of solutions for centralized reinforcement learning problem, but they are generally effective in practice and considerably more efficient to learn. Modular reinforcement learning is closely related to another variant of classical reinforcement learning: *hierarchical reinforcement learning*. Hierarchical reinforcement learning primarily focuses on temporal decomposition of goals (Barto & Mahadevan, 2003). In contrast, modular reinforcement learning focuses on problems that are naturally decomposed in terms of multiple concurrent sub-goals. While hierarchical reinforcement learning may also provide a useful modeling framework for narrative-centered tutorial planning, an investigation of hierarchical models is left for future work.

A modular reinforcement learning task is formally defined in terms of  $N$  independent MDPs  $M = \{M_i\}_1^N$ , where  $M_i = (S_i, A_i, P_i, R_i)$ .<sup>2</sup> The state space of the composite task is defined as the union of the state sub-spaces for each individual MDP:  $S = S_1 \cup S_2 \cup \dots \cup S_N$ . In many formulations of modular reinforcement learning, each MDP  $M_i$  shares an action set  $A$ . However, as will be discussed in Section 3.2, this constraint is relaxed in the current work due to the heterogeneous nature of narrative-centered tutorial planning. This has limited theoretical impact, as solutions to modular reinforcement learning problems are typically approximations, and therefore lack guarantees of optimality. Therefore, the action set for the composite agent is given by the union of the action subsets for each independent MDP:  $A = A_1 \cup A_2 \cup \dots \cup A_N$ . Each agent  $M_i$  also has its own reward model  $R_i$ . Traditionally, work on modular reinforcement learning has assumed that individual agents' rewards are comparable (Karlsson, 1997; Sprague & Ballard, 2003), and they can be summed to calculate a composite reward signal for the joint agent:  $R(s_t, a, s_{t+1}) = \sum_{i \in N} R_i(s_t, a, s_{t+1})$ . However, Bhat and colleagues (2006) argue that individual agents' rewards are often not comparable in practice. We also share the same position in this dissertation. Therefore, solutions to modular reinforcement learning in this context do not necessarily optimize a single composite reward signal; solutions are composed of multiple independent, concurrent sub-solutions that must

---

<sup>2</sup> Notation for the formal description of modular reinforcement learning is borrowed from (Bhat, Isbell, & Mateas, 2006) and (Sprague & Ballard, 2003).

be arbitrated among in a principled and effective manner. Therefore, the solution to a modular reinforcement learning problem is a set of  $N$  concurrent policies:  $\pi^* = \{\pi_i^*\}_1^N$ , where  $\pi_i^*$  is the “optimal” policy for a single constituent MDP  $M_i$ . Any circumstance where two policies  $\pi_i$  and  $\pi_j$ , with  $i \neq j$ , recommend different actions in the same state requires an arbitration procedure to be employed to select an appropriate action. It should be noted that the policy obtained for each constituent MDP may not be theoretically guaranteed to be optimal. Guarantees of optimality are predicated on the following assumptions: state representations are fully Markovian, the environment does not change from learning-time to run-time, and the decision-making agent selects all future actions according to an optimal policy. However, the use of approximate states, off-line learning, and policy arbitration procedures in the current work violates these assumptions. Regardless, modular reinforcement learning is expected to yield “good” policies that are effective in practice.

Modular reinforcement learning provides several benefits—such as learning efficiency and representational naturalness—at the expense of modest reductions in overall optimality (Karlsson, 1997). These tradeoffs naturally align with the requirements of narrative-centered tutorial planning. For example, in cases where narrative-centered tutorial planning models are induced directly from student interaction data, training data is generally limited in quantity. Consequently, learning efficiency is critical. As a second example, modular reinforcement learning formalizations model multiple concurrent goals as distinct rewards for each sub-task. This is contrasted with centralized reinforcement learning formalizations that attempt to combine metrics for multiple goals in a single reward function. Given the multi-goal nature of narrative-centered tutorial planning, this representational quality of modular reinforcement learning is compelling for system designers. The next section describes how narrative-centered tutorial planning can be modeled in terms of concurrent Markov decision processes, laying the foundations for modular reinforcement learning-based solutions to narrative-centered tutorial planning.

### 3.2 Narrative-Centered Tutorial Planning with Concurrent Markov Decision Processes

Narrative-centered tutorial planning involves dynamically adapting interactive stories in order to promote effective and engaging learning experiences that are tailored to individual students. This is achieved by strategically manipulating story events as they unfold in a narrative-centered learning environment. A narrative-centered tutorial planner should be capable of performing a wide range of actions in order to modify an interactive narrative's plot, adjust its discourse, or augment the beliefs and intentions of a student user. Each opportunity to augment an interactive narrative experience is an instance of an adaptable event sequence. The actions that are performed by a narrative-centered tutorial planner are narrative adaptations.

**Definition.** An *adaptable event sequence* (AES) encapsulates a series of one or more related story events that can unfold in multiple forms within an interactive narrative system. The different forms of an AES may have distinct impacts on the interactive narrative, but the forms can be interchanged with one another without affecting the narrative's coherence. An AES may involve inserting, re-ordering, augmenting, or removing story events from a "canonical" narrative sequence. Additionally, an AES can occur one or multiple times during the interactive narrative.

**Definition.** A *narrative adaptation* is a concrete sequence of events that represents a particular manifestation of an AES. Each narrative adaptation constitutes an action taken by a narrative-centered tutorial planner. Narrative adaptations are typically selected based on a user's actions, preferences, or needs.

To illustrate the concept of an AES, consider the following example: an event sequence that occurs after a student asks a virtual character about her backstory. The virtual character may respond in several ways. She may provide a detailed explanation about her

backstory, she may respond suspiciously and reveal only a few details, or she may refuse to respond at all. In this scenario, each of the character's three possible responses is an alternate manifestation of a "character backstory" AES. The system's choice about how the virtual character should respond is an example of a narrative adaptation. When a narrative-centered tutorial planner performs a narrative adaptation, its "action" is a decision about which manifestation of the AES should occur.

Using the concept of an AES, narrative-centered tutorial planning can be cast as a collection of sequential decision-making problems about AESs impacting students' narrative-centered learning experiences. In effect, each AES is a sub-problem in the overarching narrative-centered tutorial planning task. The planner's success in performing narrative adaptations is demonstrated by students' learning outcomes, students' levels of engagement, and the quality of students' narrative experiences. These success criteria are effectively delayed rewards that can guide the decision-making model. This conceptualization—a collection of sequential decision-making tasks with delayed rewards—points to concurrent Markov decision processes as a natural approach for modeling narrative-centered tutorial planning. Figure 3 illustrates how narrative-centered tutorial planning can be modeled with such a formalization.

Modular reinforcement learning techniques address several key challenges that arise in narrative-centered tutorial planning. First, the approach enables planners to be directly induced from data collected during student interactions with a narrative-centered learning environment. Narrative-centered tutorial planning is a novel and ill-understood computational problem. Employing machine learning-based techniques presents an opportunity to obtain models empirically, which is a promising approach for devising effective systems. However, collecting log data from student interactions with narrative-centered learning environments is costly and time consuming. Even under the best circumstances, the volume of training data that can be collected from human users is orders of magnitude smaller than the simulation-generated corpora used by many reinforcement

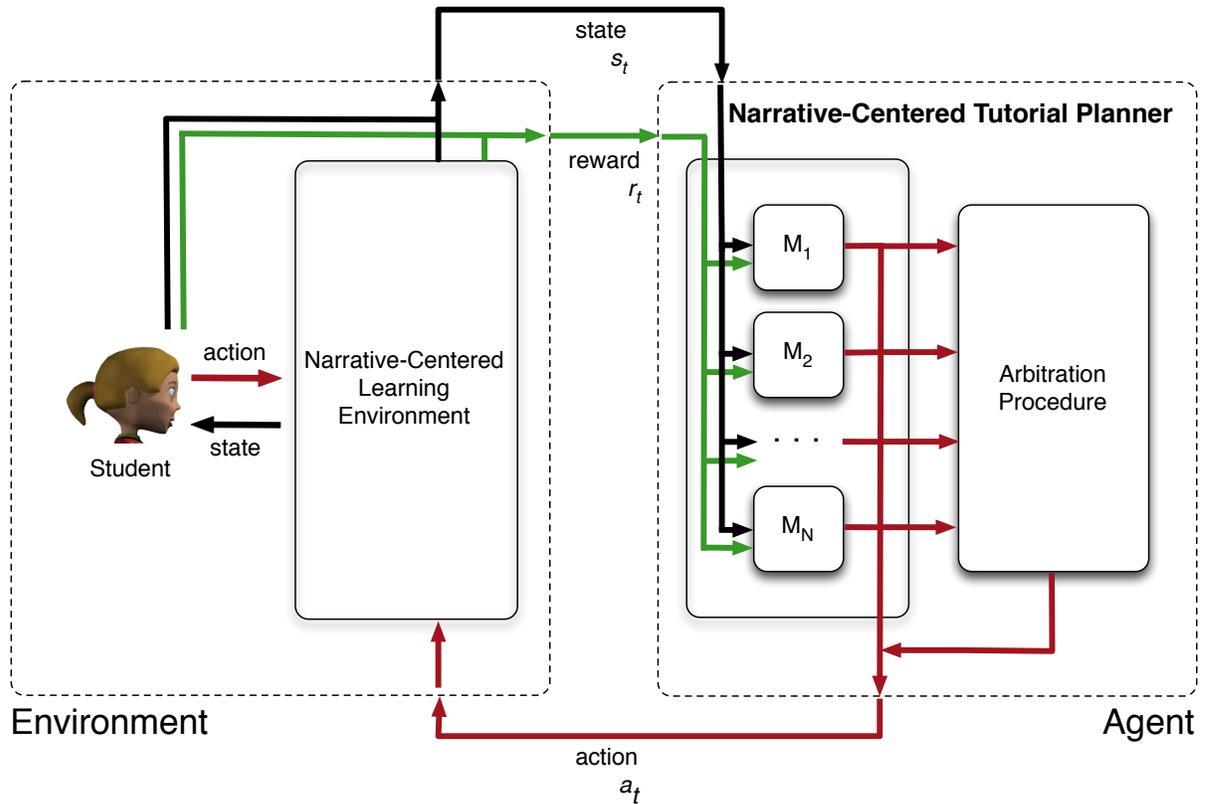


Figure 3. Agent-environment architecture for narrative-centered tutorial planner based on modular reinforcement learning.

learning systems. Therefore, it is critical to select algorithms that can efficiently solve Markov decision processes from limited data. By decomposing narrative-centered tutorial planning into multiple concurrent, independent sub-problems with restricted state representations, it is possible to learn a composite narrative-centered tutorial planner from moderately sized corpora of student log data. A critical facet of this approach is an assumption that the planner's constituent Markov decision processes are independent. In other words, we assume that narrative adaptations performed for one AES do not affect the environment states, or state transitions, of distinct MDPs representing other AESs. By adopting this assumption, we avoid the need to model all of the AESs jointly, substantially

reducing the planning task's complexity.

A second challenge of narrative-centered tutorial planning is its inherently multi-faceted nature. A narrative-centered tutorial planner is expected to simultaneously optimize student learning and problem-solving performance within the context of a coherent and engaging narrative. While these goals should be complementary in well-designed narrative-centered learning environments, they do not necessarily demand identical courses of action by a planner. Consider the following different types of policies. A narrative-centered tutorial planner may perform narrative adaptations that promote content learning during the initial stages of an interactive narrative experience, thereby providing the student with the requisite knowledge to advance through the problem-solving scenario. An example would be directing a virtual character to initiate a conversation about educational subject matter. If a student encounters an impasse, and therefore risks becoming bored, the planner may instead select engagement-enhancing adaptations. An example engagement-enhancing adaptation could be a prompt that encourages the student to take a short break from problem solving and explore an unvisited area of the virtual story world. In some ways, this narrative adaptation conflicts with content learning goals, at least in the short term, yet it is appropriate given the students' potential affective state. During the final stages of an interactive narrative experience, the planner may seek to de-emphasize content learning and provide problem-solving scaffolding that ensures the student arrives at a coherent resolution to the narrative scenario.

Each of these examples pertains to a distinct goal: content learning, engagement, and problem-solving performance, respectively. Concurrent Markov decision processes are a natural formalism for decomposing narrative-centered tutorial planning tasks based on AESs, enabling each sequence to be associated with its own sub-goals. By decomposing a narrative-centered tutorial planning task in this manner, concurrent Markov decision processes provide a mechanism for inducing narrative adaptation policies that collectively account for multiple goals.

In addition to multiple goals, narrative-centered tutorial planners should accommodate a diverse range of narrative adaptations. The intelligent narrative technologies

community has devised several methods for dynamically augmenting interactive narrative experiences (Rowe, Shores, Mott, & Lester, 2010b), and this range of narrative adaptations should be at the disposal of narrative-centered tutorial planners. Each AES that is supported by a narrative-centered tutorial planner corresponds to a well-defined set of narrative adaptations that can be executed in the virtual story world. Furthermore, each AES merits consideration of distinct state features, and the decisions contribute toward a distinct set of goals. By modeling each AES as a distinct, independent MDP—including a tailored state representation, action set, and reward signal—it is possible to accommodate a wide range of possible narrative adaptations while avoiding the curse of dimensionality associated with centralized decision-making formalizations.

A fourth challenge inherent in narrative-centered tutorial planning is the real-time performance requirements of interactive narrative systems. The solution to a set of concurrent Markov decision processes is a set of policies that map environment states to agent actions. In narrative-centered tutorial planning, this corresponds to a direct mapping between task environment states and narrative adaptation decisions. The obtained policies can be implemented as lookup tables (for small state spaces) or approximation functions (for complex state spaces). This enables planners to make narrative adaptation decisions quickly; in most cases, narrative adaptation decisions can be performed in constant time in this framework. By inducing a distinct policy (or a fixed set of policies) for each AES, the run-time performance of an AES-based narrative-centered tutorial planner meets the requirements of real-time interactive narrative systems.

Modeling narrative-centered tutorial planning as a modular reinforcement learning-problem involves four representational considerations. First, the overall task must be decomposed, which involves identifying the set of AESs and goals for the narrative-centered tutorial planner. Each AES corresponds to a sub-problem in modular reinforcement learning; in our framework, each AES is modeled as a distinct, independent MDP. Second, techniques for devising appropriate state representations for each sub-problem must be identified. This includes identifying categories of state features, as well as utilizing principled techniques for

feature selection. Third, the set of concrete event sequences that comprise each type of narrative adaptation must be formally specified. These sequences constitute the action sets for each MDP. Fourth, the task environment for each MDP must be precisely characterized, including when narrative adaptations are performed, sources of uncertainty in the environment's dynamics, and the duration of narrative experiences. The remainder of this section describes each of these four representational issues, and provides intuitions for operationalizing each step.

**Task decomposition.** A narrative-centered tutorial planner should have access to a wide range of narrative adaptations in order to achieve multiple parallel goals, such as enhancing student problem solving, student learning, and student engagement. Decomposing the overall task of narrative-centered tutorial planning involves two stages: 1) identifying a set of AESs that encompass a wide range of narrative adaptations in the narrative-centered learning environment, and 2) associating a set of quantifiable goals with each AES. Explicitly defining AESs and associated goals enables a system designer to begin formulating the MDPs that model the overall narrative-centered tutorial planning problem.

The first step of task decomposition is identifying a set of AESs that are supported by the story environment. Identifying the AESs for an interactive narrative is inherently a creative process, and it is generally the domain of the environment's designer. In some ways, AESs are similar to branch points in story graphs; AESs are key points in an interactive narrative's plot that can unfold in several possible ways, each distinctly impacting users' story experiences. However, the independence assumptions in our framework—each AES is independent of all other AESs—pose restrictions on AESs' properties, a notable distinction from traditional conceptualizations of branch points. While there is not a single “correct” process for defining the AESs in an interactive narrative, there are several criteria that one can use in identifying opportunities for AESs: 1) a designer can envision the event, or sequence of events, unfolding in one of several well-defined ways, such as variations in a character's behavior or the player's abilities, without harming the interactive narrative's

coherence; 2) the conditions for the event's occurrence arise one or multiple times in the interactive narrative, such as speaking with a particular virtual character or entering a building; 3) the event is likely to substantially impact students' experiences in the narrative-centered learning environment; and 4) different variations of the event may prove better or worse for a student depending on his individual characteristics. These are not the only criteria one can use, but they have been utilized to identify the concrete AESs that we examine later in this dissertation.

As mentioned previously, a single AES may occur one or multiple times during a student's narrative-centered learning experience. For example, consider a "prompting" AES where a virtual character chooses whether to prompt a student to record important information that she recently encountered. The student is likely to encounter useful information several times during the narrative. Correspondingly, the "prompting" AES could occur several times, and each time the student could record the new findings. Alternatively, consider a "story twist" AES, where a virtual character presents a profound revelation that alters the course of the story. This type of AES only occurs once. After the revelation has been made, it cannot be revealed again. In these scenarios, an MDP is associated with the "prompting" AES, and a separate MDP is associated with the "story twist" AES. In the former case, each successive decision about whether to prompt the student is an action by the "prompting" MDP. The separate "story twist" MDP would only perform a single binary action (i.e. choosing whether to present the revelation or not) during an episode.

The second step of task decomposition is identifying a set of goals that are the focus of the narrative-centered tutorial planner. These are the objectives that drive the planner's decision making, and they guide reinforcement learning by providing a source for the MDPs' reward signals. In the case of a narrative-centered tutorial planner that seeks to enhance student learning outcomes, student engagement, and student problem-solving performance, each of these three objectives is treated as a distinct goal. After identifying a goal set, metrics for objectively quantifying progress toward each goal are formulated. For example, learning outcomes can be quantified in terms of learning gains, measured as the difference between

student scores on a content pre-test and post-test. Assessments of engagement may include in-game measures such as frequency of off-task behaviors, and out-of-game measures such as post hoc questionnaires about student presence and motivation. Problem-solving performance can be assessed in terms of the number of essential in-game problem-solving steps that a student has completed. Specific decisions about goal metrics are dependent on the interactive narrative system and instructional setting.

After determining objective metrics for each goal, a mapping between AESs and goals must be defined. This mapping is largely based on the designer's intended purposes for narrative adaptations, such as scaffolding problem solving or promoting engagement. In situations where a single AES corresponds to two (or more) goals, the AES can be represented using two (or more) MDPs with one goal each. In the prior example of the "prompting" AES, each prompt is intended to improve the student's problem solving performance, as well as his content learning gains. However, problem-solving performance and learning gains are quantified differently, and they are not easily compared. Techniques for combining these two factors into a single scalar measure would likely be arbitrary or invalid. Therefore, the "prompting" AES should be represented using two MDPs: one that optimizes problem solving performance, and one that optimizes content learning.

After developing a mapping between AESs and goals, the goal metrics must be transformed into well-defined reward signals. These are the actual quantities that will be optimized during model induction. A reward signal may simply reproduce raw values from an AES's goal metric, or it may transform the metric's value in order to adjust its range, cardinality, or distribution. Rewards may be provided at the conclusion of a narrative, or they may be distributed incrementally over the course of a decision-making episode. However, care must be taken when utilizing incremental rewards. By rewarding incremental behaviors, the learner may produce a model with a sub-optimal policy that benefits incremental goals at the expense of global objectives.

In computing MDP policies that optimize rewards, one must clearly define the conception of "optimality" employed by solution techniques. Computing narrative-centered

tutorial planning policies requires clear optimization criteria, which define how reward signals are interpreted to assess the values of different states and actions. In modeling narrative-centered tutorial planning with Markov decision processes, we seek policies that maximize total accumulated rewards during episodes. However, we also seek to promote efficient policies. To promote efficiency, we discount future rewards relative to immediate rewards. A common mechanism for encoding discount mechanisms in MDPs is the discount factor,  $\lambda$ , which we employ in this dissertation.

In summary, task decomposition proceeds as follows. First, we identify the AESs in the narrative-centered tutorial planning task, which correspond to sub-problems in modular reinforcement learning. Second, we identify planning goals for each AES, which are used to define reward models for the associated MDPs. In conceptualizing how the constituent MDPs will be solved, we will seek to compute policies that maximize the total accumulated reward obtained during episodes, where future rewards are discounted relative to proximal rewards.

**State representation.** Interactions with narrative-centered learning environments yield rich and diverse data about students' problem solving actions, narrative experiences, and individual differences. These data sources collectively comprise the task environment's state, and they can be utilized to inform narrative-centered tutorial planning. In an ideal setting, a narrative-centered tutorial planner would leverage a comprehensive range of features that fully characterizes the student's state, the virtual environment's state, and the interactive narrative-so-far. However, even if it were possible to devise such a state representation, it would likely be computationally prohibitive. As features are added to the state, the dimensionality of the machine learning task increases, resulting in exponential growth in the problem's complexity. Furthermore, excessively detailed state representations may yield fine-grained policies that are overfitted to their training sets. Therefore, a balance must be struck between using too many features and using too few features in the state representation. A narrative-centered tutorial planner must distinguish between situations that demand

alternate courses of action, but it must consider few enough features to ensure that the machine learning task is tractable given constraints in training data, time, and memory.

Rather than use a comprehensive collection of features, approximate states can be employed (Singh, Kearns, Litman, & Walker, 1999). Approximate states abstract the “true” state space by using a limited number of carefully chosen state features. A distinct state representation can be crafted for each MDP in the modular reinforcement learning framework, using only those state features that are most relevant to the associated AES. States are represented using vectors of feature values that are deterministically calculated from outside user data (e.g., survey data) and interactive narrative event logs. The features are typically transformed into coarse-grained, discrete scalars, which further reduces the complexity of the problem space. Fine-grained distinctions in a state feature’s range may have little impact on a narrative-centered tutorial planning policy, but they can substantially impact the size of the reinforcement learning problem space.

Previous work on drama management has used state representations that consist of abstract sequences of transpired plot points (Mott & Lester, 2006a; Nelson, et al., 2006; Yu & Riedl, 2012). This representation provides a view of the user’s plot experience, but it omits several data sources that could potentially inform narrative adaptation decisions. Work on intelligent tutoring systems has investigated a broad range of features in devising pedagogical models (Chi, Jordan, VanLehn, & Hall, 2008; Tetreault & Litman, 2006). However, interactions with most intelligent tutoring systems are more constrained than those afforded by narrative-centered learning environments, and therefore they support narrower sets of possible state features.

Three categories of state features are considered in the proposed dissertation on narrative-centered tutorial planning (Figure 4). The first category is *narrative state features*. These features resemble the representations utilized in previous work on drama management (Mott & Lester, 2006a; Nelson, et al., 2006). These features indicate which plot points have been completed at particular stages of a narrative-centered learning experience, as well as the

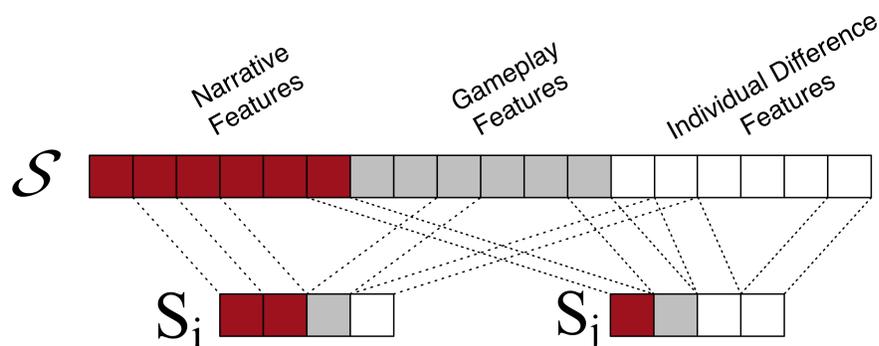


Figure 4. Feature selection for narrative-centered tutorial planner states.

relative order of the completed plot points. They can also highlight which plot points are available to be accomplished next, as well as high-level information about narrative structure. The second category is *gameplay features*, which encompass summary statistics about how students have interacted with the narrative-centered learning environment, including their problem solving and exploratory behaviors. Example features from this category include statistics regarding students' conversations with virtual characters (dialogue branches traversed, total amount of time in conversations), interactions with virtual objects (when did the student first pick up a particular object), navigational characteristics (time spent in specific locations, efficiency in traversing the virtual environment), and interactions with curricular and narrative content (time spent reading virtual books about educational subject number).

A third class of state features is *individual difference features*, which summarize distinctive characteristics of each student. Previous work has illustrated that individual differences can help predict how students will interact with a narrative-centered learning environment, as well as resulting learning and engagement outcomes (Rowe, et al., 2010a). By extension, leveraging information about students' individual differences provides a useful source of data to inform narrative-centered tutorial planning decisions. Examples of individual difference features include the level of students' prior curricular knowledge, their

previous game-playing experience, personality characteristics, gender, and self-efficacy for the subject matter.

Manually selecting features to be included in the state representation can be done in an ad hoc manner, or based on prior empirical and theoretical work. However, manual feature selection can be challenging, as there are few guarantees that selected features will be effective for driving narrative adaptation decisions for a given AES. Furthermore, manual feature selection can grow unwieldy when authoring states for a large number of modular MDPs. Fortunately automated techniques that search for effective state representations are also available (Chi, et al., 2008). These techniques typically draw upon a large pool of pre-specified state features for potential inclusion. The search process begins by learning an optimal policy for a state representation that uses only a single feature. After obtaining the policy, several metrics can be calculated that quantify the policy's value, such as expected cumulative reward (ECR), confidence intervals for the ECR, or the policy's hedge (Chi, et al., 2008; Tetreault, Bohus, & Litman, 2007). These metrics can be used to compare alternate policies, and therefore they can be used to guide a systematic search process. The search procedure proceeds iteratively, systematically swapping and adding features to the state representation, then learning and comparing policies using the previously described metrics. The state representation that is ultimately produced is the one with the greatest combination of expected cumulative reward and reliability, while also conforming to constraints in state size and search resources.

An additional note should be made about state representations in narrative-centered tutorial planning. So far, this section has described the state representation issue as involving a complex, fully observable environment. However, in reality the task environment is partially observable. For example, the planner does not have direct access to the student's internal beliefs, knowledge, or intentions. A deliberate choice has been made in this work to avoid representing narrative-centered tutorial planning as a partially observable MDP. This decision is in the interest of using training data that has been collected from actual student interactions with a narrative-centered learning environment. Explicitly representing partial

observability in narrative-centered tutorial planning would demand significantly greater quantities of training data, significantly simpler state representations, or both. Instead, the task environment in the proposed dissertation is treated as fully observable. Hidden state features are considered as additional sources of uncertainty, and no explicit effort is made to model these features. Explicitly modeling the partially observable nature of narrative-centered tutorial planning is left for future work.

**Action specification.** In order to define the narrative adaptations that are supported by a narrative-centered learning environment, the set of concrete variations for each AES must be explicitly enumerated. In our framework, each AES is associated with an MDP. The concrete variations for a particular AES are represented by the action set for the corresponding MDP. Actions in an MDP are directly mapped to narrative adaptations, or sequences of concrete events within the narrative-centered learning environment. Consider the earlier example of the “prompting” AES. In this example, a virtual character makes a binary choice between whether to prompt the student about completing a particular problem-solving step, or not prompt the student. The corresponding MDP’s action set contains two actions: 1) prompt and 2) don’t prompt. When the planner considers the “prompting” AES, it chooses a narrative adaptation from this binary action set.

Distinct AESs rarely share actions or action sets. Collections of AESs and narrative adaptations are ideally broad and heterogeneous, providing the narrative-centered tutorial planner with a diverse range of options for tailoring an interactive narrative. In circumstances where there is overlap between multiple AESs, it may even be appropriate to combine them.

Several categories of narrative adaptations have been investigated by the intelligent narrative technologies community, including plot adaptations, discourse adaptations, and user tailoring (Rowe, Shores, Mott, & Lester, 2010b). Plot adaptations are structural manipulations of a fully or partially ordered sequence of story events for an interactive narrative experience. Plot adaptations can be realized in many forms, such as inserting and removing story events or modifying virtual character states. Discourse adaptations involve

augmenting the presentation of an interactive narrative experience. They can be implemented through variations in cinematography or through variations in event presentation sequences, such as analepsis and foreshadowing. User tailoring describes customized guidance, support, and feedback that do not significantly augment the plot structure or discourse of the interactive narrative experience. User tailoring may surface as individualized cognitive and affective guidance messages, or discreet coaching from a virtual character about self-regulated learning processes. Each of these categories corresponds to a distinct family of possible action types within a narrative-centered learning environment.

One notable limitation of our narrative-centered tutorial planning framework is that certain types of narrative adaptations are not well-supported; AESs should not include actions that prevent the occurrence of later plot points in the interactive narrative. Riedl and colleagues (2003) refer to these actions as *exceptional events*. They are events with effects that threaten the preconditions of future narrative events. The present narrative-centered tutorial planning framework does not guarantee that induced policies will avoid coherence-disrupting narrative adaptations, particularly while the model is being machine learned. In principle one would expect the framework to avoid such actions because they likely lead to low rewards. However, presenting incoherent narratives to students in educational settings is ill-advised. We recommend that the narrative adaptations for each AES observe the following property: they possess no effect that can invalidate the preconditions of an essential plot point occurring later in the interactive narrative. This simplifying assumption—avoiding narrative adaptations that cause *exceptional events*—provides guarantees of coherent narrative-centered learning experiences, while accommodating dynamic tailoring to individual students.

**Task Environment.** The task environment for an MDP describes the conditions under which actions are performed. In the proposed narrative-centered tutorial planning framework, the task environment involves state transitions that directly or indirectly result from narrative adaptations. Each instance of an AES introduces a *decision point*, which involves a planning

decision about which narrative adaptation to perform. Decision points can be sequenced in two primary ways. They can occur at regular time intervals, such as every few seconds or minutes. Depending on the interval's duration, regularly occurring decision points may result in a majority of AESs producing no action (i.e. the planner does not need to continuously manipulate the narrative). Alternatively, decision points can be triggered based on specific conditions within the environment, such as talking to a particular character, entering a specific location, or completing a critical plot point. Decision points with well-defined preconditions yield simpler reinforcement learning problems, but they constrain the number of circumstances in which a narrative-centered tutorial planner can act. The current work focuses on irregular decision points that are triggered by environmental conditions.

In narrative-centered tutorial planning, the MDP's state transition model describes how the task environment changes its state between successive decision points. Specifically, at decision point  $t$ , the MDP for the  $i$ -th AES, denoted  $M_i$ , is in state  $s_{i,t}$ . At this point, the optimal policy for  $M_i$  recommends an action  $a_i$ , which is consequently executed in the narrative-centered learning environment. As a consequence of the narrative adaptation,  $M_i$  transitions to a successor state in accordance with the state transition model's probabilities. The subsequent MDP state  $s_{i,t+1}$  is calculated at the *next decision point* for the  $i$ -th AES. This means that the successor state  $s_{i,t+1}$  does not necessarily describe the conditions that immediately follow the narrative adaptation. In fact, a considerable duration may pass before the next decision point. If no subsequent decision point occurs, the final observed state of the narrative-centered learning experience is used to determine  $s_{i,t+1}$ . In this framework, several events may occur in between successive decision points for a particular AES. These events are likely to result in incremental changes to the environment state. However, these incremental changes are not directly considered by the policy for  $M_i$ . Events that occur in between decision points are treated as sources of uncertainty that impact the state transitions calculated from successive decision points.

Task environments in narrative-centered tutorial planning are marked by inherent uncertainty. This uncertainty arises from several sources. First, student actions are difficult to

predict, but they are key drivers of interactive narrative experiences. Student actions are motivated by plans, intentions, beliefs, and affective states that are not directly accessible by the system, and they introduce considerable modeling challenges of their own. Second, approximate states capture only a constrained subset of the environment state by definition. Factors that are not explicitly represented in an MDP's state representation may still impact the transitions observed in the task environment. This uncertainty can be explicitly modeled by probability values in the MDP's state transition model, assuming that the model can represent probability distributions resembling the real environment's dynamics.

A final consideration in formulating task environments for narrative-centered tutorial planning is whether the decision-making task is episodic or continuing. Most narrative-centered learning experiences are episodic. They observe some form of narrative structure with a beginning, middle, and end. However, some genres of narrative continue without end, such as soap operas that continually introduce new characters and story twists. In modular reinforcement learning, these infinite narratives can be modeled as continuing tasks. Continuing narrative experiences may offer promise for educational applications, where a student can continue exploring curricular content and problem-solving activities in a manner that is contextualized within an ongoing narrative. By removing a terminal state from a narrative-centered learning experience, artificial constraints for time-on-task may also be removed. However, generating endless interactive narrative experiences and problem-solving activities introduce their own distinct challenges. The framework presented in this dissertation focuses on the more common type of episodic narrative-centered learning experiences for narrative-centered tutorial planning.

This section has described the operational details of representing narrative-centered tutorial planning with concurrent Markov decision processes. After formulating the problem representation, the next step is to induce an action policy for the composite narrative-centered tutorial planner. The following section discusses data-driven techniques for obtaining effective narrative-centered tutorial planning policies.

### 3.3 Inducing Narrative-Centered Tutorial Planning Policies

In our concurrent Markov decision process framework, the overall narrative-centered tutorial planner is an abstraction for several independent models that make decisions about individual AESs. Creating the overall narrative-centered tutorial planning model proceeds in two stages: 1) inducing the individual policies for each AES, and 2) employing arbitration techniques that merge complimentary and competing policies. This section describes reinforcement learning techniques for inducing policies from actual student data. Additionally, techniques and considerations for policy arbitration are presented.

**Policy induction.** In reinforcement learning, a distinction is made between *model-free* learning techniques and *model-based* learning techniques (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). Model-free techniques assume that no explicit environment model is available for obtaining an optimal MDP policy. Model-free techniques learn from experience, simultaneously exploring a task environment and pursuing policies that accumulate maximal reward. They include approaches such as Monte Carlo methods and temporal-difference techniques (Sutton & Barto, 1998). Model-based techniques, such as dynamic programming methods, require explicit environment models that describe state transition probabilities and expected reward values. Although there are typically no explicit environment models available in narrative-centered tutorial planning tasks, the models can be approximated using data obtained during dedicated corpus collection stages, which is an application of certainty equivalent methods (Kaelbling, Littman, & Moore, 1996).

Model-based techniques and model-free techniques also give rise to another distinction in reinforcement learning procedures: *on-line* learning and *off-line* learning. On-line learning involves data collection, policy learning, and policy execution stages that are performed concurrently. The learner must balance between *exploring* the state space and *exploiting* its obtained knowledge in order to maximize accumulated reward (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). As a narrative-centered tutorial planner learns increasingly effective policies, this knowledge can be used to collect data from more

promising parts of the problem space and exploit policies that yield greater rewards. On-line reinforcement learning models are generally trained using simulation data, as opposed to data collected from human users. Work on declarative optimization-based drama management has investigated on-line reinforcement learning methods using simulated users and simulated interactive narrative environments (Nelson, et al., 2006; Roberts, et al., 2006).

On-line reinforcement learning techniques offer several benefits when training models from simulation-generated data. However, notable challenges arise when training a narrative-centered tutorial planner from human user data. During large data collections involving human users, it is generally necessary to have multiple participants concurrently interact with the narrative-centered learning environment on separate computers. Each of these narrative-centered learning interactions is independent; an on-line learner on one computer is unaware of user interactions taking place on other computers. In order to perform aggregate on-line reinforcement learning in this setting, coordination mechanisms would be necessary to share learned state-action values among independent narrative-centered tutorial planner instances. This coordination introduces synchronization challenges, such as determining which individual learner has the most up-to-date value estimate for a state-action pair. Alternatively, coordination issues could be avoided by performing an entire data collection on a single computer. However, restricting a data collection to one computer would be logistically prohibitive. The quantity of training data that is required by reinforcement learning would result in data collections with enormously long durations. As a consequence, on-line reinforcement learning methods introduce significant practical difficulties when training a narrative-centered tutorial planning policy directly from human user data.

Off-line learning refers to a class of procedures that separate the data-collection and model-operation stages. When training a narrative-centered tutorial planning policy from human user data, off-line learning requires that users interact with a version of the narrative-centered learning environment that is specifically designed for collecting a training corpus. This narrative-centered learning environment should be identical to the final system—

including the set of AESs that are supported—with the exception of the policies used to drive narrative adaptation decisions. The data collection system should perform narrative adaptations in a manner that is exploratory, such as a random policy, rather than a manner that seeks to maximize accumulated reward. This enables a broad sampling of the state space, producing data that can be used to calculate an approximate environment model. After the corpus data has been collected, model-based reinforcement learning techniques can be employed to induce a set of “optimal” narrative adaptation policies. The resulting policies are then implemented in a new, deployable version of the narrative-centered learning environment. In this method, the cohort of students used to generate data for reinforcement learning are distinct from the students who interact with the eventual machine-learned narrative-centered tutorial planner. This approach is a form of certainty equivalent learning (Kaelbling, Littman, & Moore, 1996).

The environment model for each AES can be calculated from a corpus of student log data as follows. The state transition model for an MDP  $M_i$  refers to the probability  $P(s_{i,t+1} | a_t, s_{i,t})$ . This probability can be calculated by counting the total number of instances in the corpus where  $M_i$  took action  $a_i$  while in state  $s_{i,t}$  and then transitioned to state  $s_{i,t+1}$ , divided by the total number of instances where  $M_i$  took action  $a_i$  while in state  $s_{i,t}$ . By repeating this calculation for all possible combinations of states and actions, the full state transition model can be obtained for a single AES. The reward model  $R_i$  refers to the expected reward  $E(r_{i,t+1} | s_{i,t}, a_t, s_{i,t+1})$ . This value can be calculated by determining the average observed reward for  $M_i$  when the action  $a_i$  was taken in state  $s_{i,t}$  and resulted in a transition to state  $s_{i,t+1}$ . After calculating an approximate environment model and reward model from the collected corpus data, each MDP  $M_i$  is fully specified. At this stage, dynamic programming algorithms such as value iteration or policy iteration can be employed to determine the values of each MDP state  $s_i$  (Sutton & Barto, 1998). This procedure can be repeated for each MDP  $M_i$  in order to calculate an approximately optimal policy for each independent AES.

Off-line learning offers several benefits for training a narrative-centered tutorial planner from human user data. First, it avoids the coordination and logistical challenges that arise during on-line learning with human users. There is no coordination necessary during off-line learning's corpus collection stage, because each planner instance performs narrative adaptations in an exploratory manner. Second, by uniformly sampling the space of possible narrative adaptation policies, off-line learning procedures enable investigations of several candidate state representations and reward models. Comparisons among alternate state representations can be performed using the policy metrics described in Section 3.2.

However, off-line learning approaches do possess limitations. They collect data inefficiently, sampling from both promising and unpromising areas of the state space. This limitation mandates that simple state representations be selected when formulating MDPs, thereby ensuring adequate data is available for approximating the environment model. Furthermore, the framework's implicit assumption that all AESs are independent may reduce the effectiveness of policies obtained through off-line learning. If the independence assumption is violated, then the corpus collection environment (where narrative adaptation policies are computed) will feature different environment dynamics than the deployment environment (where the narrative adaptation policies operate). In the corpus collection version of the environment, the observed state transitions are contingent on every other narrative adaptation policy acting in an exploratory manner. However, in the deployment environment every other narrative adaptation policy selects actions based on a locally optimal policy. If the individual AESs are not independent, then a policy that was effective in the corpus collection environment may not be as effective in the deployment environment. Despite these limitations, it is believed that the benefits associated with inducing narrative-centered tutorial planning policies from human user data will outweigh the costs associated with off-line learning in our concurrent Markov decision process framework.

**Policy arbitration.** In the proposed narrative-centered tutorial planning framework, each AES is generally associated with a single MDP policy. However, there may be circumstances

where multiple decision points are triggered simultaneously or multiple MDPs are associated with a single AES. The latter scenario occurs when an AES is associated with multiple goals, and it has been decomposed into several distinct MDPs. In this situation, the narrative-centered tutorial planner may receive multiple action recommendations for a given decision point. If all of the action recommendations agree, no arbitration is necessary. However, if the recommended actions differ, arbitration techniques must be employed in order to choose a single action for the composite planner.

Previous work on modular reinforcement learning has investigated arbitration procedures that exclusively consider the state-action values obtained for each module. Example arbitration procedures of this type include *greatest mass* and *nearest neighbor* strategies (Karlsson, 1997). These arbitration procedures are referred to as domain-independent, and they require that rewards for all MDPs be comparable. Later in this dissertation, we will investigate an implemented narrative-centered tutorial planner that utilizes a domain-independent arbitration procedure, because each AES will be mapped to a single reward function computed from students' normalized learning gains. However, Bhat and colleagues (2006) have argued that domain-independent arbitration procedures are inadequate for settings where modules' rewards are not comparable. In some cases, narrative-centered tutorial planning may require models involving several distinct and incomparable rewards. In these circumstances, the arbitrator must be imbued with domain-specific knowledge about the relations and tradeoffs between individual modules. This arbitrator employs knowledge that is unavailable at the individual module level, and the arbitrator uses the knowledge to dictate selections among competing policies.

In multi-goal narrative-centered tutorial planning, domain-specific arbitration procedures involve consideration of the overall instructional priorities of the narrative-centered learning environment, as well as the associated AESs. By considering the instructional context of a particular decision point, the arbitrator can weight one policy's recommendations differently than another's. For example, consider a narrative-centered learning environment that is designed to serve as a study aid in a classroom setting. This type

of narrative-centered learning environment should prioritize policies that seek to increase learning gains. In contrast, a narrative-centered learning environment that is designed for informal educational settings may prioritize policies that preserve student engagement and increase motivation. These latter narrative-centered learning environments are not primarily designed to maximize content knowledge per say, but they are designed to foster further educational exploration by enhancing student interest in the subject matter.

### 3.4 Theoretical Analysis

Decomposing narrative-centered tutorial planning into multiple concurrent Markov decision processes has a number of theoretical benefits over centralized techniques. By factoring the narrative-centered tutorial planning task into independent AESs, we take advantage of assumed independence relationships to reduce the planning task’s computational demands. Specifically, this formalization enables reductions in the quantity of training data, as well as worst case training time, required to induce narrative-centered tutorial planning policies. In this section, we provide theoretical arguments for these benefits, and describe the conditions under which they occur.

In order to demonstrate how our concurrent MDP framework for inducing narrative-centered tutorial planners reduces training corpus size requirements compared to a centralized MDP framework, consider the following argument. A centralized narrative-centered tutorial planner is modeled as a single centralized Markov decision process  $M = (S, A, P, R)$ . The state set contains a finite number of states  $|S|$ , and each state is represented as a vector of discrete features. The MDP’s action set is finite and includes  $|A|$  distinct actions. In order to employ off-line solution techniques, such as value iteration, it is necessary to estimate a state transition model  $P$  that encodes  $|S| \times |A| \times |S| = |S|^2 |A|$  different state transitions.<sup>3</sup> Assume it requires  $k$  observations to estimate each probability value in the state transition model. Assuming the corpus collection uniformly samples the state-transition

---

<sup>3</sup> The same number of values must also be calculated for the estimated reward model  $R$ . A parallel argument holds for computing reward models for centralized vs. decomposed planners.

space, a centralized narrative-centered tutorial planner would require at least  $Q_C = k|S|^2|A|$  observations during model training.

Next, consider a comparable decomposed narrative-centered tutorial planner, which factors the planning task in terms of  $N$  distinct AESs, each modeled by an independent Markov decision process  $M_i = (S_i, A_i, P_i, R_i)$ . According to our framework, every MDP  $M_i$  possesses a tailored state representation, such that the features in  $S_i$  are a subset of the features in  $S$ . Assume that  $|S_i| \leq |S| - d \forall i = 1, \dots, N$  and  $d$  is non-negative. Based on our framework's assumption that the decomposed planner's AESs are independent, we assume that  $A_1 \cap A_2 \cap \dots \cap A_N = \{\}$  and  $A_i \subset A \forall i = 1, \dots, N$ . Further, we assume  $|A| = |A_1| + |A_2| + \dots + |A_N|$ .

**Theorem.** If  $d < 2|S|$ , then the total number of observations required for the decomposed narrative-centered tutorial planner  $Q_D$  will be  $2k|A||S|d - k|A|d^2$  fewer than the number required for the centralized narrative-centered tutorial planner  $Q_C$ .

*Proof.* By the same argument presented about the size of  $Q_C$ , estimating the state-transition models in the decomposed narrative-centered tutorial planner requires the following number of training observations:  $k|S_i|^2|A_i| \forall i = 1, \dots, N$ . Or,

$$\begin{aligned} Q_D &= \sum_1^N k|S_i|^2|A_i| = k \sum_1^N |S_i|^2|A_i| \\ &\leq k \sum_1^N (|S| - d)^2|A_i| \\ &= k(|S| - d)^2 \sum_1^N |A_i| \end{aligned}$$

Recall that  $|A| = \sum_1^N |A_i|$ . Therefore,

$$\begin{aligned} Q_D &\leq k(|S| - d)^2|A| \\ &= k|S|^2|A| - (2k|A||S|d - k|A|d^2) \\ &= Q_C - (2k|A||S|d - k|A|d^2) \end{aligned}$$

Note that when  $2k|A||S|d - k|A|d^2 > 0$ , then  $Q_D < Q_C$ . This inequality holds when

$$2|S|d - d^2 > 0$$

$$2|S| - d > 0$$

$$2|S| > d$$

Thus, when  $2|S| > d$ , then  $Q_C > Q_D$  and the difference between these two quantities is  $2k|A||S|d - k|A|d^2$ .

□

In order to illustrate this result with an example, consider a concrete narrative-centered tutorial planning task with a 6-bit state vector, such that  $|S| = 64$ , and  $|A| = 16$ . A centralized narrative-centered tutorial planner that encodes the problem with a single MDP would require  $(64 \times 64 \times 16) \times k = 65,536k$  training observations to approximate the state transition model. Alternatively, consider a decomposed narrative-centered tutorial planner that divides the problem into two AESs, where the associated MDPs possess 4-bit state vectors, such that  $|S| = 16$ , and action sets containing 8 narrative adaptations. Assuming uniform sampling, the decomposed narrative-centered tutorial planner would only require enough data to produce  $2 \times (16 \times 16 \times 8) \times k = 4,096k$  observations of narrative adaptations, a reduction of approximately 94% in required training corpus size. We should note that these findings are predicated on idealized conditions, such as uniform sampling during corpus collection and a fully factorable composite narrative-centered tutorial planning task.

Decomposing a narrative-centered tutorial planning task into independent AESs can also reduce the time required to induce a model under similar conditions. We sketch an argument for this benefit by examining the time complexity of a dynamic programming technique, value iteration, for solving MDPs. The argument for alternate model-based solution techniques, such as policy iteration, is similar.

Value iteration is an iterative algorithm for solving MDPs. The time complexity for a single iteration of value iteration is  $O(|A||S|^2)$ , where  $|A|$  is the number of actions in the action set, and  $|S|$  is the number of states (Littman, Dean, & Kaelbling, 1995). The worst case

number of iterations required by value iteration increases as the discount rate  $\gamma$  approaches 1. As Littman and colleagues explain, "...value iteration can take a number of iterations proportional to  $1/(1-\gamma)\log(1/(1-\gamma))$  in the worst case" (1995). However, for fixed values of  $\gamma$ , value iteration is guaranteed to find an optimal policy in polynomial time.

In a decomposed narrative-centered tutorial planner in our framework, each constituent MDP possesses equal or fewer effective states, and fewer actions, than an equivalent centralized narrative-centered tutorial planner. However, the decomposed narrative-centered tutorial planner involves computing solutions for multiple MDPs, at least one for each AES, whereas the centralized planner requires solving only a single MDP. Reductions in worst-case training time for the decomposed narrative-centered tutorial planner are dependent on the particular formulations of the constituent MDPs. If each state representation  $S_i$  possesses a subset of features of the state representation for the centralized planner  $S$ , then  $|S_i| < |S|$ . Consequently, the time complexity for a single iteration of value iteration is reduced, because the quadratic term possesses an outside effect compared to the time required to solve multiple MDPs. For example, if  $|S_i| = |S|/N$ , then the worst case time complexity for a single iteration of value iteration across all of the MDPs would be on the order of  $N|A|(|S|/N)^2 = |A||S|^2/N$ , a reduction by a factor of  $N$ . We assume that the number of iterations required by value iteration remains constant across both scenarios, as it is proportional to the discount factor and not the number of actions.

However, if each state representation  $S_i$  is identical to the state representation for the centralized planner  $S$ , such that  $S_i = S$ , but each action set  $A_i$  is smaller than the action set for the centralized planner  $A$ , such that  $|A_i| < |A|$ , there is no reduction in time complexity for the decomposed narrative-centered tutorial planner. The complexity of a single iteration of value iteration is linear in the number of actions, and any reduction in action set sizes for constituent MDPs is mitigated by the need to solve multiple MDPs. For example, if the decomposed narrative-centered tutorial planner is modeled in terms of  $N$  MDPs, and each constituent action set  $A_i$  includes  $|A|/N$  actions, the worst case time complexity for a single iteration of value iteration across all of the MDPs is on the order of  $N(|A|/N)|S|^2 = |A||S|^2$ .

In summary, decomposing a narrative-centered tutorial planner into multiple constituent MDPs can enable substantial reductions in training corpus size requirements through feature selection with state representations. Reductions in time complexity are also dependent on reduced state representations for each constituent MDP, because reductions in action set sizes are counteracted by the need to solve multiple MDPs. The time complexity benefits are theoretically appealing, but in practice training time is rarely a bottleneck. In general, limited availability of student interaction data requires highly constrained MDP definitions (i.e., small state and action sets), producing planning problems that can be solved in short periods of time (i.e., minutes) on most personal computers.

This chapter has described a general framework for modeling narrative-centered tutorial planning with concurrent Markov decision processes. A consolidated description of the overall modeling and induction procedure is shown in Table 1. Efforts have been made to include operational recommendations in this chapter, but specific details regarding the framework's implementation in a narrative-centered learning environment have been intentionally omitted. The remainder of this dissertation will examine a concrete implementation of the framework with a test bed narrative-centered learning environment, CRYSTAL ISLAND. We begin the investigation by describing CRYSTAL ISLAND in the next chapter.

Table 1. Procedure for modeling a narrative-centered tutorial planner with concurrent Markov decision processes.

1. Define set of adaptable event sequences for narrative-centered learning environment.
  - a. Specify range of narrative adaptations for each adaptable event sequence.
  - b. Define preconditions for adaptable event sequences' decision points.
2. Implement adaptable event sequences in narrative-centered learning environment. Narrative adaptations should be performed according to exploratory (random & uniform) policy.
3. Perform data collection with human subjects interacting with narrative-centered learning environment.
  - a. Record detailed logs of student actions in environment.
  - b. Utilize pre/in-game/post measures to assess students' experiences.
4. Identify goals (and associated metrics) for the narrative-centered tutorial planner. Goal metrics are calculated using measures from step 3b.
5. Map each adaptable event sequence to one or more goals.
6. Transform the goal metrics into well-defined reward functions.
7. Define a Markov decision process  $M_i$  for each combination of narrative adaptation and goal.
  - a. Specify the action set  $A_i$  using the range of event variations from step 1.
  - b. Define the reward function  $R_i$  using the reward functions from step 6.
8. For each MDP  $M_i$ , define a state estimator  $S_i$  using manual or automated techniques. For discussion of automated state-representation search procedures, see (Chi, et al., 2008).
9. For each MDP  $M_i$ , define an approximate environment model using corpus data from step 3, and the partial MDP definitions from steps 7-8.
  - a. 
$$P(s_{i,t+1} | a_i, s_{i,t}) = \frac{\text{Total \# of transitions from state } s_{i,t} \text{ and action } a_i \text{ to state } s_{i,t+1}}{\text{Total \# of observations involving state } s_{i,t} \text{ and action } a_i}$$
  - b. 
$$E(r_{i,t+1} | s_{i,t}, a_i, s_{i,t+1}) = \frac{\text{Sum of rewards from state } s_{i,t} \text{ and action } a_i \text{ transitioning to state } s_{i,t+1}}{\text{Total \# of transitions from state } s_{i,t} \text{ and action } a_i \text{ to state } s_{i,t+1}}$$
10. Use dynamic programming method (i.e. value iteration, policy iteration) to obtain an "optimal" policy  $\pi_i^*$ .
11. Specify domain-specific arbitration procedure to select among competing policies.
12. For each narrative adaptation in narrative-centered learning environment, replace exploratory policy (step 2) with learned "optimal" policy  $\pi_i^*$  and arbitration procedure from steps 10-11.

## CHAPTER 4

### CRYSTAL ISLAND

In order to investigate how students learn from narrative-centered learning environments, we developed CRYSTAL ISLAND, a test bed narrative-centered learning environment for inquiry-based learning of middle school microbiology. Although narrative-centered learning environments vary widely in presentation, gameplay, and instructional design, CRYSTAL ISLAND typifies open-ended environments that require students to seek out, interpret, and synthesize embedded science information in order to apply newly acquired knowledge toward solving a story-centric problem-solving task. In terms of traditional game genre classifications, CRYSTAL ISLAND most closely resembles 3D single-player adventure games. The player freely explores a virtual storyworld to gather clues and converse with non-player characters while solving an overarching mystery. Initial versions of CRYSTAL ISLAND have been non-adaptive, but they provide a foundation for developing the adaptive narrative-centered learning environment that is the subject of this dissertation.

This chapter describes core aspects of CRYSTAL ISLAND's science narrative and gameplay elements. Additionally, findings from several empirical analyses of student learning and problem solving with a non-adaptive version of CRYSTAL ISLAND are presented. The findings draw upon a single dataset of middle school student interaction logs, and they are primarily descriptive; the results offer an empirical account of how students learn from CRYSTAL ISLAND, interact with CRYSTAL ISLAND, and problem solve in CRYSTAL ISLAND. Furthermore, we examine the impacts of students' individual differences, such as gender and game-playing experience, on students' learning and interaction profiles. Developing an empirically based understanding of students' narrative-centered learning experiences establishes a foundation for designing adaptive systems such as narrative-centered tutorial planners. In combination, these investigations illustrate the challenges and opportunities

presented by narrative-centered learning environments, which point toward the need for devising narrative-centered tutorial planners.

#### **4.1 CRYSTAL ISLAND Narrative-Centered Learning Environment**

In its fourth major iteration, CRYSTAL ISLAND is a narrative-centered learning environment built on Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The curriculum underlying CRYSTAL ISLAND's mystery narrative is derived from the North Carolina state standard course of study for eighth-grade microbiology. The environment is designed as a supplement to classroom instruction, and it blends elements of both inquiry learning and direct instruction. Due to its capacity to teach and engage students, CRYSTAL ISLAND also offers promise as a tool for preparation for future learning (Schwartz & Bransford, 1998). Over the past several years, CRYSTAL ISLAND has served as a platform for investigating a range of AI technologies for dynamically supporting students' learning experiences. This includes work on combined narrative and tutorial planning (Mott & Lester, 2006b), narrative director agents (Mott & Lester, 2006a), student modeling (Lee, Mott, & Lester, 2010; Mott, Lee, & Lester, 2006; Rowe & Lester, 2010), archetype-driven models of character dialogue (Rowe, Ha, & Lester, 2008), empathetic character behavior models (Robison, McQuiggan, & Lester, 2009), and affect recognition models (McQuiggan, Mott, & Lester, 2008). These AI technologies have primarily been implemented in branches of the core CRYSTAL ISLAND environment; they have not yet been re-incorporated into the main system.

CRYSTAL ISLAND's premise involves a mysterious illness that is afflicting a research team stationed on a remote island. The student plays the role of a visitor who recently arrived on the island in order to see her sick father. However, the student gets drawn into a mission to save the entire research team from the spreading outbreak. The student explores the research camp from a first-person viewpoint and manipulates virtual objects, converses with characters, and uses lab equipment and other resources. In order to solve the mystery,

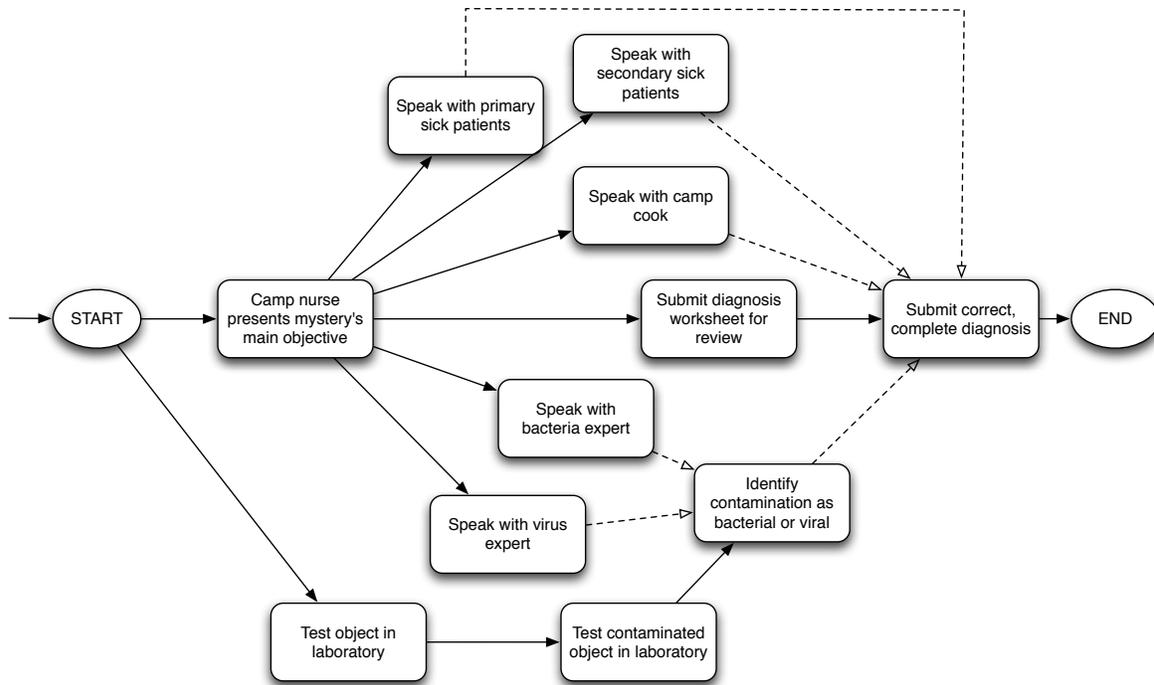


Figure 5. Partial order among sub-goals in the CRYSTAL ISLAND narrative scenario.

students complete a series of partially ordered sub-goals that uncover details about the spreading infection.

The narrative's 11 sub-goals are presented in Figure 5. In the figure, square boxes represent the eleven sub-goals to be completed. Solid arrows indicate precondition relationships between sub-goals. Dotted arrows indicate soft precondition relationships, where completion of the source node's sub-goal contributes to completion of the destination node's sub-goal, but their relative ordering is not strictly enforced.

In order to illustrate a typical interaction with CRYSTAL ISLAND, consider the following scenario. Upon arriving at the research camp, the student observes a group of small buildings a short distance from the camp entrance. The student approaches the first building,



Figure 6. Conversation with CRYSTAL ISLAND's camp nurse.

an infirmary, where several sick patients and a camp nurse are located. The student initiates a conversation with the nurse by approaching her and clicking the mouse. The nurse explains that an unidentified illness is spreading through the camp and asks for the student's help to diagnose the disease (Figure 6). She advises the student to use an in-game diagnosis worksheet in order to record findings, hypotheses, and a final diagnosis (Figure 7). This worksheet is designed to scaffold the student's problem solving process, as well as provide a space for the student to offload any findings gathered about the illness. The conversation with the nurse takes place through a combination of multimodal character dialogue—spoken language, gesture, facial expression, and text—and student dialogue menu selections. All

Patients' Symptoms			
Symptom 1	Cough	Symptom 2	Fever
Symptom 3	No Entry	Symptom 4	No Entry

Test Results			
	I have tested:	The lab results showed:	
Object 1	Water	Not contaminated	
Object 2	Bananas	Not contaminated	
Object 3	No Entry	No Entry	
Object 4	No Entry	No Entry	
Object 5	No Entry	No Entry	
Object 6	No Entry	No Entry	

Hypotheses			
	Likelihood of this diagnosis	Because	
Anthrax	No Entry	No Entry	
Botulism	Unlikely	Characteristics don't match	
Ebola	No Entry	No Entry	
Influenza	No Entry	No Entry	
Salmonellosis	No Entry	No Entry	
Smallpox	Very Unlikely	Characteristics don't match	

Final Diagnosis			
Disease	No Entry	Treatment	No Entry
Source	No Entry	Infection Type	No Entry

Figure 7. Diagnosis worksheet for recording findings in CRYSTAL ISLAND.

character dialogue is provided by voice actors and follows a deterministic branching structure.

After speaking with the nurse, the student has several options for investigating the illness. Inside the infirmary, the student can talk to sick patients lying on medical cots. Clues about the team members' symptoms and recent eating habits can be discussed and recorded in the diagnosis worksheet or as free-form notes using an in-game smartphone. Alternatively, the student can move to the camp's dining hall to speak with the camp cook. The cook describes the types of food that the team has recently been eating and provides clues about which items warrant closer investigation. In addition to learning about the sick team members, the student has multiple options for gathering information about different disease-

causing agents. For example, the student can walk to the camp's living quarters where she will encounter a pair of virtual scientists willing to answer questions about viruses and bacteria, respectively. The student can also learn more about pathogens by viewing posters hanging inside of the camp's buildings (Figure 8) or reading books located in a small library. In this way, the student can gather information about relevant microbiology concepts using resources that are presented in multiple formats.

After a conversation with a virtual character, the student may receive a phone call from the camp nurse (using an in-game smartphone) and be asked to answer a short series of multiple-choice questions about microbiology concepts. The nurse prefaces the questions by explaining that she is trying to determine the student's progress in solving the mystery. The questions are designed to assess what material the student has retained from the character interactions.

Beyond gathering information from virtual scientists and other instructional resources, the student can conduct tests on food objects using the laboratory's testing equipment (Figure 9). The student encounters food items in the dining hall and laboratory, and she can test the items for pathogenic contaminants at any point during the learning interaction. For each test, the student must specify the type of test she wishes to conduct and select a justification for that test. A limited number of tests are allocated to the student at the start of the scenario, but additional tests can be earned by answering microbiology multiple-choice questions. Therefore, if a student squanders her available tests by using a haphazard problem-solving strategy, she must demonstrate her understanding of microbiology concepts in order to continue advancing the story.

After running several tests, the student discovers that the sick team members have been consuming contaminated food. Upon arriving at this finding, the student is instructed to see the lab technician, Elise, for a closer look. The screen momentarily fades to black to indicate elapsing time, and Elise returns with an image of the contaminated specimen, which



Figure 8. Interior of the camp infirmary in CRYSTAL ISLAND.

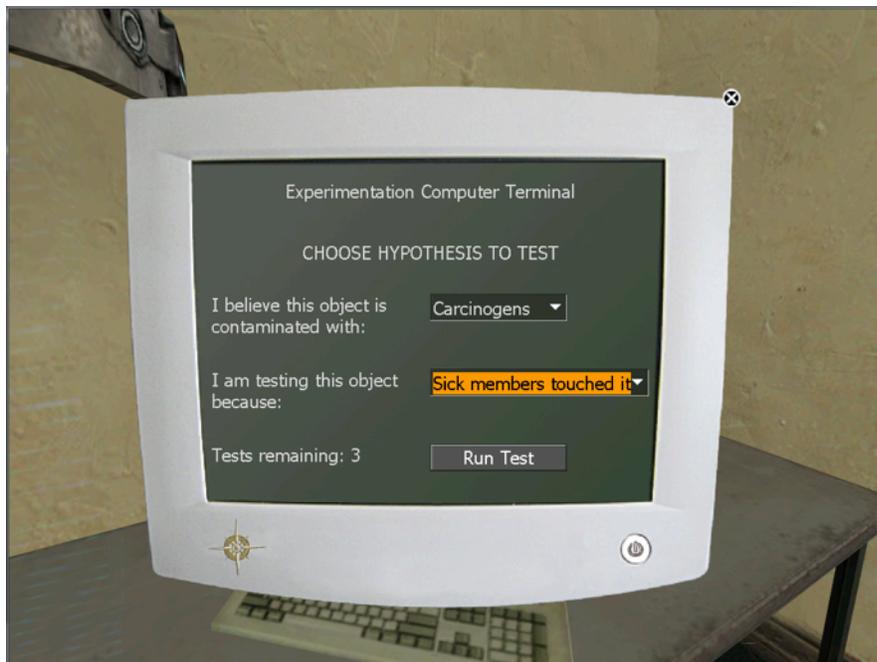


Figure 9. Testing device in the camp's laboratory.

she explains was taken using a microscope. At this point, the student is presented with a labeling exercise where she must specify the identity and parts of a pathogenic contaminant. After successfully completing this activity, the student can use the camp's books and posters in order to investigate diseases that might have produced the sick team members' symptoms. The student enters her findings and hypotheses into the diagnosis worksheet. Once she has narrowed down a diagnosis and recommended treatment, as well as entered them in her worksheet, the student returns to the infirmary in order to report her findings to the camp nurse. If the student's diagnosis is incorrect, the nurse identifies the error and recommends that the student keep working. If the student has correctly diagnosed the illness and specified an appropriate treatment, the mystery is solved.

## **4.2 Human Subject Study**

An experiment involving human participants was conducted in fall 2009 with the eighth grade population of a North Carolina middle school. The primary goal of the experiment was to investigate the impact of different scaffolding techniques on learning and engagement in the CRYSTAL ISLAND narrative-centered learning environment. However, no condition effects were observed for either learning or engagement. The findings reported in this chapter come from a secondary analysis of the data, which considers the experiment's conditions as a whole.

**Participants.** A total of 153 eighth grade students ranging in age from 12 to 15 ( $M = 13.3$ ,  $SD = 0.48$ ) interacted with the CRYSTAL ISLAND environment during the study. Data from 16 of the participants was eliminated due to incompleteness or prior experience with an earlier version of CRYSTAL ISLAND. Among the remaining students, 77 were male and 60 were female. Approximately 3% of the participants were American Indian or Alaska Native, 2% were Asian, 32% were African American, 13% were Hispanic or Latino, and 50% were White. The participants had not yet been exposed to the microbiology curriculum unit of the

North Carolina state standard course of study in their regular classes, and therefore had minimal prior experience with CRYSTAL ISLAND's microbiology content.

**Measures.** Several subjective and objective measures were used to assess student learning, problem solving and engagement in CRYSTAL ISLAND. In particular, assessing engagement in narrative-centered learning environments poses notable challenges, and it requires a multi-faceted approach. This includes using subjective and objective measures, as well as online and post hoc measures, that assess multiple dimensions of narrative-centered learning experiences. For example, presence measures are designed to assess the fidelity and authenticity of an interactive virtual environment (Witmer & Singer, 1998). Measures of situational interest can assess how appealing a narrative-centered learning experience is to students (Schraw, 1997), and they encapsulate both game-related and narrative-related contributors. A student's efficiency in advancing a plot, as well as avoidance of gaming the system behaviors, are partial on-line indicators of how engaged the student is in a narrative-centered learning scenario. The work presented in this chapter focuses on situational interest, presence, and a researcher-generated *final game score* measure that quantifies students' narrative efficiency and deliberative problem-solving behaviors. While additional measures would be necessary to achieve a comprehensive view of engagement, the current selection of instruments provides a principled first step toward assessing the multiple components of engagement in narrative-centered learning.

One week prior to interacting with CRYSTAL ISLAND, students completed an online demographic survey, game-playing experience questionnaire, self-efficacy for science measure, and CRYSTAL ISLAND curriculum test. The game-playing experience questionnaire asks students to self-report about their game-playing habits and perceived game-playing skill (Appendix A). The self-efficacy for science learning questionnaire includes eight Likert items using a five-point scale (Appendix B). The measure is adapted from (Nietfeld, Cao, & Osborne, 2006), which utilizes a portion of the measure presented in (Britner & Pajares, 2006). The curriculum test consists of 16 multiple-choice questions created by an

interdisciplinary team of researchers (Appendix C). The test consists of sixteen questions assessing students' knowledge of pathogens, select diseases, and the scientific method.

Post-experiment materials were completed immediately following the CRYSTAL ISLAND intervention. Included in these materials were the same curriculum test used in the pre-experiment, a variation of the Perceived Interest Questionnaire (Schraw, 1997), and the Presence Questionnaire (Witmer & Singer, 1998). The interest scale is adapted from (Schraw, 1997), and consists of ten Likert items measuring students' situational interest related to CRYSTAL ISLAND (Appendix D). The Presence Questionnaire (PQ) is a standard measure containing several subscales, including involvement/control, naturalism of experience and quality of interface (Witmer & Singer, 1998). Presence describes a user's sense of "being there" when interacting with a mediated environment (Schubert, Friedmann, & Regenbrecht, 1999). It is related to the sensation of being transported into a story, which is an important contributor to the engaging quality of narratives. Example items from the Presence Questionnaire include the following: "How compelling was your sense of moving around inside the virtual environment," and "How much did your experiences in the virtual environment seem consistent with your real-world experiences?"

In addition to pre- and post-experiment questionnaires, students' *in-game problem solving* performance and *final game score* were considered. In-game problem solving is measured by counting the number of sub-goals that students completed in CRYSTAL ISLAND among the narrative's 11 possible sub-goals. Not all of the sub-goals are mandatory for solving the mystery, but completing all 11 sub-goals is an important component of effective problem solving in the environment. The in-game problem solving metric treats all goals equivalently; none of the goals are weighted differently than others.

Final game score is a numerical value calculated by the CRYSTAL ISLAND software to assess students' progress and efficiency in completing the science mystery. Students could view their scores in the upper left corner of their screens throughout their interactions. The final game score consisted of a weighted sum of gameplay sub-scores, and incorporated time

Table 2. Point values for calculating final game score.

Action	Points (pts)
<b>Overall Mystery Solution</b>	
Correct Solution	500 pts
Solution Efficiency	(7500 / elapsed time) pts
Incorrect Solution Attempt	-100 pts
<b>In-game Quiz Questions</b>	
First Attempt Correct	25 pts
Second Attempt Correct	10 pts
Second Attempt Incorrect	-10 pts
<b>Object Contaminant Testing</b>	
Test Milk for Pathogens	200 pts
Incorrect Object	-10 pts
Incorrect Contaminant	-25 pts
<b>Character Interactions</b>	
Talk to Kim	(25 / elapsed time) pts
Talk to Teresa	(50 / elapsed time) pts
Talk to Ford	(125 / elapsed time) pts
Talk to Robert	(125 / elapsed time) pts
Talk to Quentin	(125 / elapsed time) pts
<b>Pathogen Labeling Activities</b>	
Correct Answer	10 pts
Incorrect Answer	-10 pts
<b>Total Maximum Points</b>	<b>≈ 1665 pts</b>

taken (in seconds) to accomplish important narrative goals, students' ability to demonstrate microbiology content knowledge, and evidence of careful hypothesis formulation. Students were penalized for any attempt to "game the system" by repeatedly submitting incorrect diagnoses to the camp nurse or guessing on content knowledge quizzes. Gaming the system is a common form of disengagement in intelligent tutoring systems. Details of final game score's calculation are shown in Table 2.

In the current work, final game score is used to capture aspects of student behavior associated with engagement in the narrative-centered learning environment. Final game score quantifies the degree to which students exercised learning behaviors that are associated with

engaged problem solving, and avoided behaviors associated with guessing and gaming the system. While the measure does incorporate non-engagement related factors, such as demonstration of microbiology knowledge, these factors are lightly weighted in the overall calculation. It is not proposed that final game score measures engagement directly, but rather that it is correlated with engagement and provides insight into students' learning experiences beyond what is available from post hoc questionnaires. As an objective online measure characterizing students' gameplay characteristics in the CRYSTAL ISLAND environment, final game score provides a tool for distinguishing students who were disengaged in the scenario (as evidenced by inefficient goal completion or repeated attempts to guess the solution) and students who were engaged in the problem-solving task.

**Procedure.** Participants entered the experiment room having completed the majority of pre-test materials one week prior to the intervention. Students were provided general details about the CRYSTAL ISLAND mystery and game controls during an introductory presentation by a researcher. After the presentation, students completed the remaining pre-experiment materials and received several CRYSTAL ISLAND supplementary documents. These materials consisted of a CRYSTAL ISLAND backstory and task description, a character handout, a map of the island, and an explanation of the game's controls (Appendix E).

Participants were given 60 minutes to work on solving the mystery. Immediately after solving CRYSTAL ISLAND's science mystery, or after 60 minutes of interaction, participants completed the post-experiment questionnaires. Completion of post-experiment materials took no longer than 30 minutes for participants. In total, sessions lasted up to 120 minutes.

### **4.3 Student Learning, Problem-Solving and Engagement Findings**

This section summarizes findings about student learning, problem solving, and engagement from the CRYSTAL ISLAND study<sup>4</sup>. An investigation of student learning found that on average,

---

<sup>4</sup> The statistical analyses were conducted collaboratively with Lucy Shores.

Table 3. Regression predicting students' microbiology post-test scores.

Predictor	B	SE	$\beta$
Pre-Test	.46**	.09	.33**
Presence	.03*	.01	.15*
Final Game Score	.01**	.00	.31**
$R^2$		.33	

Note: \*\* -  $p < .01$ ; \* -  $p < .05$

students answered 2.35 ( $SD = 2.75$ ) more questions correctly on the post-test than they did on the pre-test. Matched pairs t-tests (comparing post-test to pre-test scores) indicated that students' learning gains were significant,  $t(149) = 10.5, p < .001$ .

Examining factors believed to reflect engagement and students' understanding of the curriculum, Pearson correlations indicated significant relationships between microbiology background knowledge and presence ( $r = .17, p < .05$ ) and final score ( $r = .28, p < .01$ ). Similar relationships were found between microbiology post-test scores and presence ( $r = .29, p < .01$ ), final score ( $r = .45, p < .01$ ), and situational interest ( $r = .24, p < .01$ ). To more closely investigate the relationships between learning and engagement, additional analyses controlling for background knowledge were conducted.

A partial correlation controlling for pre-test score found significant relationships between microbiology post-test scores and two of our engagement measures, presence ( $r = .25, p < .01$ ), and final game score ( $r = .38, p < .01$ ). The same type of analysis also found a borderline significant relationship between situational interest and post-test score ( $r = .15, p < 0.1$ ). Offering further evidence for a connection between learning and engagement in CRYSTAL ISLAND, a linear regression indicated that microbiology background knowledge, presence, and final game score were all significant predictors of performance on the microbiology post-test, and the model as a whole was significant ( $R^2 = .33, F(3, 143) = 23.5, p < .001$ ; see Table 3).

Table 4. Hierarchical regression predicting students' microbiology post-test scores.

Predictor	Model 1			Model 2		
	B	SE	$\beta$	B	SE	$\beta$
Pre-Test	.61**	.10	.44**	.46**	.11	.32**
Presence				.03*	.02	.16*
Goals Completed				.54**	.13	.33**
Situational Interest				.02	.03	.05
R <sup>2</sup>		.18			.35	

Note: \*\* -  $p < .01$ ; \* -  $p < .05$

Analyses were conducted to investigate the relationships between engagement-related measures and in-game problem solving ( $M = 9.26$ ,  $SD = 1.91$ ). In these analyses, only presence and situational interest were investigated as engagement-related variables. The analyses do not consider final game score, because its calculation is based in part on how efficiently students completed a subset of CRYSTAL ISLAND's eleven goals.

Pearson correlations revealed significant relationships between number of goals completed and microbiology background knowledge ( $r = .17$ ,  $p < .05$ ), microbiology post-test scores ( $r = .47$ ,  $p < .01$ ), situational interest ( $r = .21$ ,  $p < .05$ ), and presence ( $r = .27$ ,  $p < .01$ ). Similarly, a partial correlation analysis controlling for microbiology background knowledge found significant relationships between number of goals completed and microbiology post-test performance ( $r = .40$ ,  $p < .01$ ) and presence ( $r = .24$ ,  $p < .01$ ).

To further investigate the relationship between in-game problem solving and engagement, a hierarchical regression analysis was conducted controlling for microbiology background knowledge (see Table 4). In order to predict microbiology post-test performance, microbiology pre-test score was entered into the first block while situational interest, presence, and number of goals completed were entered into the second block. Microbiology background knowledge was found to be a significant predictor of microbiology post-test performance ( $F(1, 130) = 28.6$ ,  $p < .01$ ) and was responsible for 18% of the variance.

However, the second model was also found to significantly predict microbiology post-test performance ( $F(4, 127) = 16.9, p < .01$ ). Microbiology pre-test score, number of goals completed, and presence were all significant predictors of microbiology post-test performance, and accounted for 35% of post-test score's variance. Situational interest was not observed to be a significant predictor.

The findings indicate that student engagement with the CRYSTAL ISLAND environment was associated with improved learning and problem-solving outcomes. Results showed a significant relationship between students' pre-test scores and presence, as well as between pre-test scores and final game scores. This suggests that students who demonstrated greater prior content knowledge tended to become more engaged with the narrative environment. However, all three measures for engagement—presence, situational interest, and final game score—were found to be significantly associated with post-test score, and these associations extended beyond their relationship with pre-test score. Additionally, presence and situational interest were found to be significantly associated with in-game problem-solving performance, also extending beyond their relationship with pre-test score. These findings suggest that CRYSTAL ISLAND was capable of engendering significant learning gains, and that engagement in the narrative-centered learning environment was an important contributor to student learning and problem solving. The results provide evidence that CRYSTAL ISLAND is an effective educational tool, and they reinforce the promise of devising a version of CRYSTAL ISLAND with a narrative-centered tutorial planner that enhances student learning, problem solving and engagement outcomes through personalized scaffolding.

#### **4.4 Individual Difference Findings**

An analysis of students' final in-game scores was conducted<sup>5</sup>. Students' final in-game scores averaged 353 points ( $SD = 513$ ), and ranged from -1095 points to 1317 points. This

---

<sup>5</sup> These analyses were also performed collaboratively with Lucy Shores.

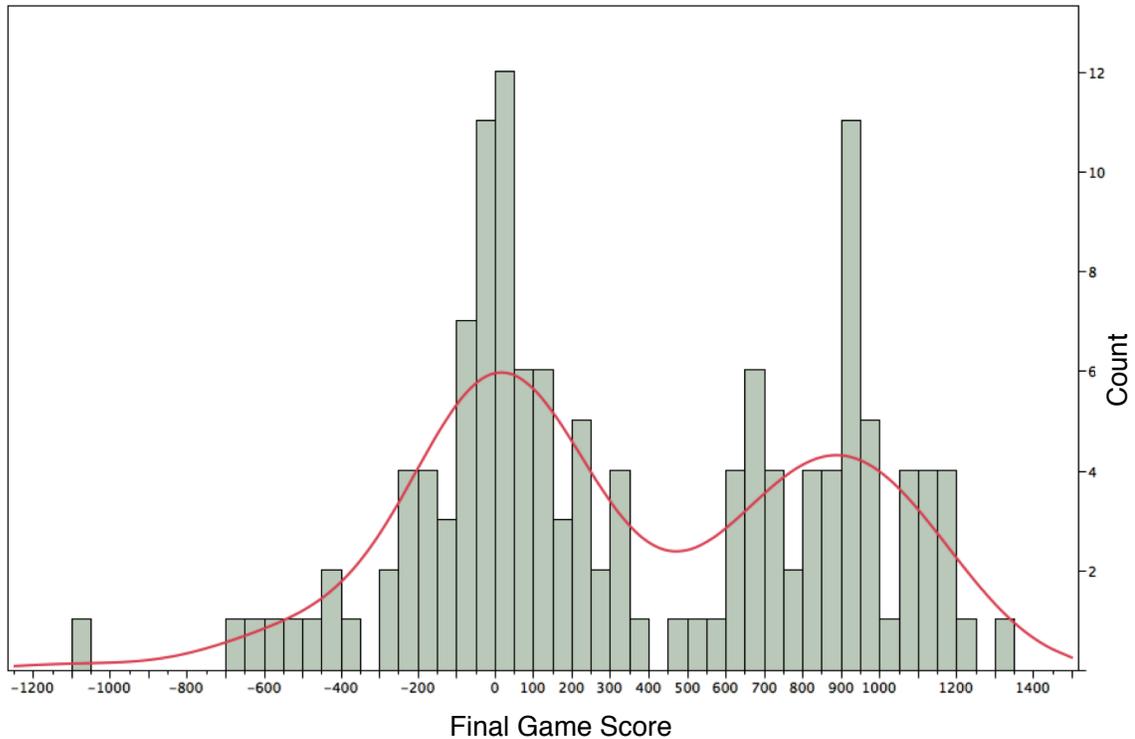


Figure 10. Distribution of students' final game scores.

analysis revealed a bimodal distribution over student scores (Figure 10); among the 137 participants considered in the analysis, 79 students scored in the lower range of -1095 to 450 points, and 58 students scored in the upper range of 451 to 1317 points. In other words, slightly more than half (58%) of the population clustered around the low end of the score range, and slightly less than half (42%) of the population was clustered near the high end. Further analysis was performed in order to investigate this polarization among gameplay scores.

The participant data was partitioned into high-scoring (greater than 450 points) and low-scoring (less than 450 points) groups. Several ANOVA tests were performed to

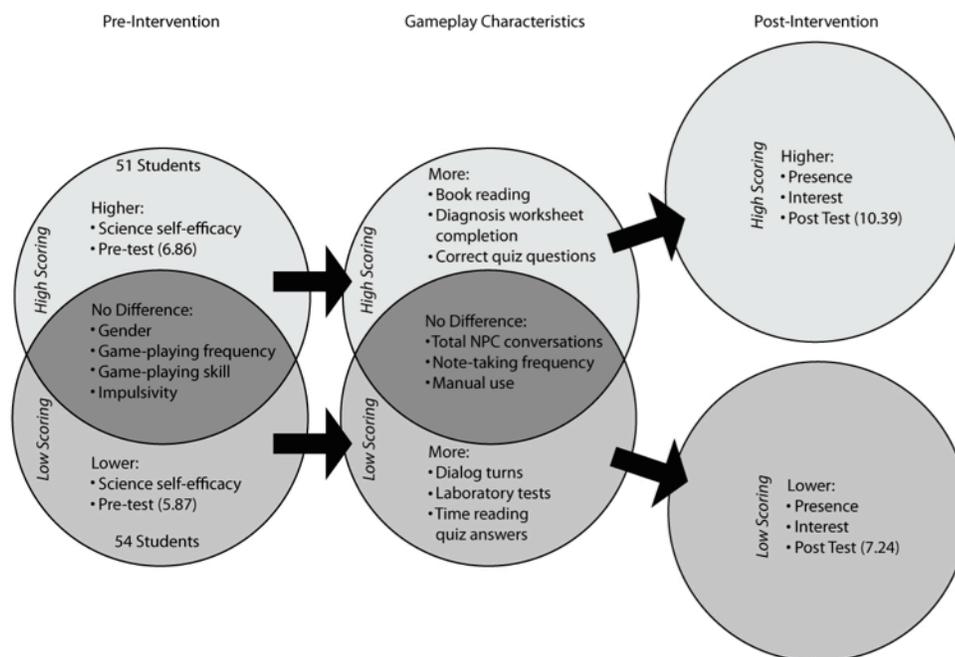


Figure 11. Individual differences prior to gameplay, in-game, and after gameplay.

investigate differences between the two clusters. A summary of the findings is shown in Figure 11. Investigating the pre-game measures revealed that the high-scoring group was significantly more self-efficacious for science,  $F(1, 100) = 16.9, p < .0001$ , and had greater microbiology background knowledge,  $F(1, 135) = 7.75, p < .01$ . Interestingly, no significant differences were found between the groups for either game-playing frequency or self-perceived game-playing skill. Furthermore, no differences were found for gender.

A series of ANOVAs were performed to examine the differences between the low- and high-scoring groups on the post-game measures. Students in the high-scoring group answered significantly more questions correctly on the microbiology content test,  $F(1, 135) = 31.6, p < .0001$ , and reported higher levels of situational interest,  $F(1, 128) = 6.99, p < .01$ , and presence,  $F(1, 135) = 22.5, p < .0001$ . Because the high-scoring group demonstrated greater microbiology background knowledge prior to the

CRYSTAL ISLAND intervention, an ANCOVA controlling for pre-test score was conducted. The analysis revealed that the high-scoring group still performed significantly better on the microbiology content post-test,  $F(1, 134) = 23.2, p < .001$ . This indicates that the high-scoring group demonstrated greater microbiology learning gains than the low-scoring group.

An analysis of the variance between the high- and low-scoring groups indicated significant differences in students' gameplay characteristics. The high-scoring group read virtual books more often than the low-scoring group,  $F(1, 135) = 15.9, p < .0001$ , and spent more time reading books,  $F(1, 135) = 3.99, p < .05$ . High-scoring students also completed their diagnosis worksheets more accurately,  $F(1, 135) = 176.1, p < .0001$ . This was true for all of the diagnosis worksheet subsections: patient symptoms,  $F(1, 135) = 19.6, p < .0001$ , laboratory test findings,  $F(1, 135) = 20.2, p < .0001$ , diagnosis hypotheses,  $F(1, 135) = 46.1, p < .0001$ , and final diagnosis,  $F(1, 135) = 277.2, p < .0001$ . Perhaps not surprisingly, high-scoring students also performed better on in-game quizzes. For the set of primary quiz questions that all students should have received during gameplay, high-scoring students answered more questions correctly on the first attempt,  $F(1, 135) = 4.61, p < .05$ , and answered marginally more questions correctly in total,  $F(1, 135) = 3.44, p = .07$ .

In contrast, low-scoring students spent their time conducting more laboratory tests than the high-scoring group,  $F(1, 135) = 17.4, p < .0001$ . Low-scoring students also conducted more erroneous tests,  $F(1, 135) = 12.4, p < .001$ , and when asked to explain their choice to conduct a particular test, more often replied with an erroneous justification,  $F(1, 135) = 10.5, p < .01$ . Low-scoring students were also more frequently involved in conversational interactions with non-player characters than high-scoring students. The low-scoring group spent more total time engaged in dialogue interactions,  $F(1, 135) = 11.2, p < .001$ , and also selected marginally more dialogue branches during conversational interactions,  $F(1, 135) = 3.62, p < .06$ . Although high-scoring students performed better on quiz questions than low-scoring students, the low-scoring group spent more time reviewing quiz question answers than the high-scoring group,  $F(1, 135) = 33.3, p < .0001$ . No differences were found between the high- and low-scoring groups for number of initiated

conversations ( $M = 22.6$ ,  $SD = 7.7$ ), microbiology field manual usage ( $M = 2.3$ ,  $SD = 1.7$ ), or note-taking frequency ( $M = 2.9$ ,  $SD = 3.8$ ).

The results of the analysis reveal several interesting differences regarding how eighth grade students interact with CRYSTAL ISLAND. A striking contrast in scores, revealing two substantially different groups of students and game-play experiences, may be observed. Students who were likely to have a greater disposition for science (i.e., those with greater microbiology background knowledge and self-efficacy for learning science) tended to achieve higher scores during the game interaction, experienced improved learning outcomes compared to students demonstrating a lesser disposition toward science, and focused on distinct types of gameplay activities. It should be noted that the score difference was not a factor of gender, perceived game-playing skill, or game-playing frequency.

Students' scores were tied to their efficiency in solving CRYSTAL ISLAND's mystery, their avoidance of gaming-the-system behaviors, and their ability to demonstrate microbiology content knowledge at select points in the scenario. From the current analysis, it is difficult to determine whether the low-scoring students were off-task. However, the findings do suggest that low-scoring students focused on different activities than high-scoring students. Low-scoring students were less efficient in achieving important goals, and demonstrated less understanding of microbiology content during gameplay. Further, these students tended to converse more extensively with non-player characters and conducted larger numbers of laboratory tests. However, many of these tests were poorly chosen and unnecessary, and the students did not appear to always use the results effectively toward solving the overarching mystery. This may be symptomatic of student difficulties in forming and following effective problem-solving strategies.

In contrast, high-achieving students appear to have been more engaged in accumulating the microbiology knowledge necessary for solving the mystery, as shown by increased time reading virtual books and correctly answering curriculum quiz questions. High-scoring students were also more effective at using their diagnosis worksheets to organize findings and record hypotheses. Both high-and low-scoring students spent relatively

equal amounts of time reading the microbiology manual and taking notes. These findings reveal important individual differences in how students interact with narrative-centered learning environment such as CRYSTAL ISLAND. The findings also demonstrate that individual differences can have significant effects on students' learning and engagement outcomes. Incorporating narrative-centered tutorial planners, which can provide tailored problem-solving support to individual students, appears to be a promising approach for assisting students who otherwise may struggle in narrative-centered learning environments.

## CHAPTER 5

### Training Corpus Collection

While empirical evidence indicates that the non-adaptive version of CRYSTAL ISLAND enhances students' content knowledge and problem-solving processes, the findings also suggest that individual differences in students' science dispositions significantly shape learning outcomes. The literature on instructional design makes clear that instructors, whether human or computer-based, should provide guidance to students for developing deeper learning skills, as these skills do not automatically materialize for many learners (Kirschner & Clark, 2006; Hmelo-Silver, Duncan, & Chinn, 2007; Jonassen, 2000). However, overt pedagogical actions are not in the spirit of narrative-centered learning, which aims to promote student engagement through immersion in evocative plots, compelling gameplay, and believable characters. Designers of narrative-centered learning environments aim to embed tutorial guidance discreetly. Instructional support should help students to learn, as well as advance interactive plots, but it should not disrupt students' experiences of narrative transportation.

Crafting effective narrative-centered tutorial planning policies presents notable challenges, largely due to the young state of the research field, and an associated lack of established design principles. Rather than entirely depend on human designers to effectively encode narrative-centered tutorial planning policies by hand, data-driven methods present opportunities to induce policies that explicitly optimize for student problem-solving outcomes within designer-specified constraints. Furthermore, data-driven methods have the potential to reveal novel policies, which may never have been proposed by a human author, through largely automated processes. Ideally, induced policies detail how narrative events can deliver instructional support discreetly while optimizing for desired outcomes, and they are gleaned from student interaction data that is *ecologically valid*. In this context, ecological

validity refers to how closely the training data correlates to student behaviors in authentic settings, such as classroom-based deployments led by teachers.

This chapter describes the process of extending CRYSTAL ISLAND with a narrative-centered tutorial planner induced from student interaction data. First, modifications to the CRYSTAL ISLAND narrative-centered learning environment are described, including a set of thirteen adaptable event sequences that were incorporated into the environment's interactive narrative. The following sub-section describes two classroom studies that were conducted to collect training data for automatically inducing narrative-centered tutorial planning policies. During the studies, students interacted with the modified version of CRYSTAL ISLAND, which performed narrative adaptations in an exploratory manner (i.e. according to a random policy) in order to systematically sample the space of potential policies. The resulting training corpus is described in the chapter's third sub-section.

## 5.1 Modified CRYSTAL ISLAND

Chapter 4 described a non-adaptive version of the CRYSTAL ISLAND narrative-centered learning environment, which provides the basis for an adaptive version imbued with a data-driven narrative-centered tutorial planner. In the non-adaptive version of CRYSTAL ISLAND, whenever a student performs an action in the virtual world, the interactive narrative responds in a fixed manner according to hand-authored rules. If the player performs the same action a second or third time, the environment responds in the same way that it did during the first instance. Plot state has minimal impact on the way narrative events unfold,<sup>6</sup> and gameplay characteristics and individual differences have no bearing on narrative events.

In order to create a modified version of CRYSTAL ISLAND that personalizes story events during run-time, it was first necessary to identify which aspects of CRYSTAL ISLAND's interactive narrative should be designated as candidates for personalization. Specifically, we

---

<sup>6</sup> A notable exception to this rule occurs at the onset of the narrative. Before the student converses with the camp nurse for the first time, all of the other non-player characters direct the student toward speaking to the nurse. The nurse introduces the problem-solving scenario, and the non-player characters do not use their typical dialogue trees until the student has received this introduction.

identified plot points in CRYSTAL ISLAND's narrative that could be transformed into adaptable event sequences, and thus unfold in several possible ways without introducing *exceptional events* later in the plotline. It should be noted that this process involved retrofitting an existing interactive narrative structure with adaptable events sequences. This demonstrates that it is not required to originally devise an interactive narrative plot with adaptable event sequences in mind, although there are likely to be benefits in taking such an approach.

In total, thirteen different adaptable event sequences were identified within CRYSTAL ISLAND's storyline, each with 2-6 possible variations in their associated narrative events. The process for identifying adaptable event sequences was driven by several factors: 1) we selected events that would embody three distinct categories of narrative adaptations: plot adaptations, discourse adaptations, and user tailoring (Rowe, Shores, Mott, & Lester, 2010b); 2) we chose narrative elements that had previously been found to significantly relate to students' learning and problem-solving outcomes (Shores, Rowe, & Lester, 2011; Rowe, McQuiggan, Robison, & Lester, 2009; Spires, Rowe, Mott, & Lester, 2011; Rowe, Shores, Mott, & Lester, 2010a); and 3) we selected events that were critical to how students received, and processed, information required to solve the mystery. We elected to incorporate thirteen distinct adaptable event sequences (rather than one or two) in order to have a broad range of mechanisms to shape and manipulate students' narrative-centered learning experiences. The thirteen adaptable event sequences, and associated decision points, are as follows:

- **Diagnosis Worksheet Feedback.** When a student submits her diagnosis worksheet to the camp nurse for review, the nurse provides immediate feedback about its correctness. If there is an error in the Final Diagnosis section of the worksheet, the nurse selects one of three possible variations in her feedback: *minimal detail*, which simply asserts that an error is present; *moderate detail*, which calls attention to the particular field that contains the error; and *maximum detail*, which indicates the field containing the error and provides a hint about how to correct it. The narrative adaptations in this AES are plot adaptations.

- **Quentin's Revelation.** When a student converses with the camp cook for a second time, the cook may optionally *deliver* or *not deliver* a plot revelation: he was also sick earlier. If he delivers the revelation, the cook's dialogue tree is modified to include an extra branch that 1) elaborates on his earlier symptoms, 2) mentions the food item that likely caused his illness, and 3) suggests that the player test the suspected food in the laboratory. This revelation expedites a major phase of the narrative: identifying the disease's transmission source. The decision point for this revelation is also triggered whenever the player speaks with the camp nurse for a second time. If the revelation is activated through the camp nurse, the nurse immediately suggests that the student should go to the Dining Hall to speak with the cook. This narrative adaptation only occurs if Bryce's Revelation has not also occurred. The narrative adaptations in this AES are plot adaptations.
- **Bryce's Revelation.** When a student converses with the sick lead scientist, Bryce, for a second time, the scientist may optionally *deliver* or *not deliver* a plot revelation: he also investigated the spreading illness prior to the student's arrival. If Bryce delivers the revelation, he explains that his notes are on a password-protected computer in an adjacent room, and he provides the password. Unlocking the computer displays an e-mail message that 1) reveals which food item tests positive for contaminants, 2) indicates whether the infection is viral or bacterial, 3) provides a complete list of the disease's symptoms, and 4) hints that the student should read about relevant diseases in the laboratory's library. This message explicitly provides most of the key information necessary for solving the mystery, with the exception of the solution. This narrative adaptation only occurs if Quentin's Revelation has not also occurred. The decision point for this adaptable event sequence is also triggered whenever the player speaks with the camp nurse for a second time. The narrative adaptations in this AES are plot adaptations.
- **Off-Task Behavior Discouragement.** When a student appears to be off-task, she may optionally *receive* or *not receive* a text message on her in-game smartphone

- encouraging her to re-focus on solving the mystery. This message is framed as coming from the camp nurse. Decision points for this adaptable event sequence occur whenever the student has not entered a building for more than seven consecutive minutes. The narrative adaptations in this AES are forms of user tailoring.
- **Record Findings Reminder.** When a student uncovers information that is relevant to diagnosing the illness, she may optionally *receive* or *not receive* a reminder to record her findings in the diagnosis worksheet. Decision points for these reminders occur immediately after a student reads a virtual book (occurring after every third book), immediately after a student runs a laboratory test (occurring after every third test), and shortly after a student converses with a sick patient about symptoms (occurring after every conversation of this type). Reminders are delivered in the form of text messages from the camp nurse on the student's in-game smartphone. The narrative adaptations in this AES are forms of user tailoring.
  - **Increase Urgency.** If the student has not solved the mystery after 20 minutes, which is typically the case, she may optionally find that all of the characters begin to reply to conversational queries with an increased sense of urgency. This adaptable event sequence is a one-time binary decision; characters may *deliver* their remaining dialogue in an urgent manner, or they may *not deliver* their remaining dialogue in an urgent manner. For example, at the onset of the narrative, the camp nurse typically greets the student with a non-urgent "Hi" or "Hello." However, if the Increase Urgency decision point has resulted in dialogue with enhanced urgency, the camp nurse may instead greet the student by saying, "The team members are getting worse! Complete the Final Diagnosis section of your worksheet, then come talk to me. Hurry!" Similarly, one of the sick patients may begin greeting the student by saying, "Uugh, I am starting to feel worse. Please hurry!" Each character possesses a distinct set of greetings for urgent and non-urgent plot variations. The narrative adaptations in this AES are discourse adaptations, because they shape how the narrative is told through the characters.

- **Knowledge Quiz.** When a student first converses with the camp's cook, virus specialist, or bacteria specialist, he may optionally *receive* or *not receive* a call from the camp nurse on his in-game smartphone. If he receives the call, the nurse explains that she would like to check on the student's progress in solving the mystery, and she proceeds to ask a series of 2-6 multiple-choice questions about CRYSTAL ISLAND's microbiology curriculum. After speaking with Ford, the virus specialist, the questions pertain to virus topics. After speaking with Robert, the bacteria specialist, the questions pertain to bacteria topics. After speaking with Quentin, the camp cook, the questions pertain to general characteristics of pathogens, mutagens, and carcinogens. Each opportunity for a "quiz" is a distinct binary decision point. In other words, a student may receive some quizzes, but not others, based on the narrative-centered tutorial planner's successive decisions. The narrative adaptations in this AES are forms of user tailoring.
- **Next Goal Prompt.** After 15 minutes of interaction, the student may optionally *receive* or *not receive* a text message from the camp nurse on his in-game smartphone that includes a prompt about an in-game goal to pursue. The prompted goal is randomly selected from the set of plot goals that have not yet been accomplished by the student. An example prompt is as follows: "Kim: You should test food items with the lab's testing equipment." The decision point for the Next Goal Prompt AES occurs repeatedly every 15 minutes. The narrative adaptations in this AES are forms of user tailoring.
- **Details of Teresa's Symptoms.** When a student converses with the sick female scientist, Teresa, she may choose to inquire about recent symptoms. Teresa will reply with one of three possible variations: *minimal detail*, which reveals only a single non-discriminatory symptom, *moderate detail*, which describes two semi-discriminatory symptoms, or *maximum detail*, which provides a comprehensive account of all symptoms. Discriminatory symptoms enable the student to deduce which infections are most likely responsible for the outbreak. This decision point occurs every time the

- student inquires about Teresa's symptoms. The narrative adaptations in this AES are plot adaptations.
- **Details of Bryce's Symptoms.** When a student converses with the sick lead scientist, Bryce, he may choose to inquire about recent symptoms. Bryce will reply with one of two possible variations: *minimal detail*, which describes one discriminatory symptom, and *moderate detail*, which describes two discriminatory symptoms. This decision point occurs every time the student inquires about Bryce's symptoms. The narrative adaptations in this AES are plot adaptations.
  - **Mystery Solution.** The solution to CRYSTAL ISLAND's mystery scenario is defined by two independent factors: the disease's identity, and the disease's transmission source. There are six distinct solutions to the mystery, which are derived from a combination of two possible diseases and three possible transmission sources. The six solutions are as follows: *influenza transmitted by contaminated eggs*, *influenza transmitted by contaminated milk*, *influenza transmitted by contaminated sandwiches*, *salmonellosis transmitted by contaminated eggs*, *salmonellosis transmitted by contaminated milk*, and *salmonellosis transmitted by contaminated sandwiches*. The decision point for the Mystery Solution AES occurs once at the onset of the interactive narrative. The narrative adaptations in this AES are plot adaptations.
  - **Reflection Prompt.** After the student solves the mystery, or the scenario's time limit expires, she may optionally *receive* or *not receive* a prompt to write about her reflections on her problem-solving methods. The prompt instructs the student to leave instructions for a future scientist who is faced with saving his fellow islanders. The prompt is accompanied by a free-entry textbox, which the student may fill with as much (or little) description as she wishes. The narrative adaptations in this AES are forms of user tailoring.
  - **Initial Lab Test Count.** Students use a virtual scanning device in the camp laboratory to test objects for disease-causing contaminants. Students are allotted a

limited number of tests at first in order to discourage problem-solving strategies that involve blindly testing objects. The number of tests made available to students is determined when students first interact with the scanning device. There are three possible variations of Initial Lab Test Count: *3 tests*, *5 tests*, and *10 tests*. If a student wishes to continue running tests after depleting his initial allotment, he must answer quiz questions using the in-game smartphone to “earn” more. The earning process is not affected by the Initial Lab Test Count AES. The narrative adaptations in this AES are plot adaptations.

The adaptable event sequences described above are the only modifications added to the CRYSTAL ISLAND narrative-centered learning environment. None of the narrative adaptations introduce exceptional events that cause downstream coherence conflicts in the interactive narrative’s plot. In cases where a narrative adaptation does affect downstream events, the events have been parameterized to support alignment with all possible plot variations. For example, sick scientists report symptoms that are consistent with narrative adaptations of the Mystery Solution AES. Similarly, the scanning device produces test results that are consistent with narrative adaptations of the Mystery Solution AES.

## **5.2 Corpus Collection Studies**

After modifying the CRYSTAL ISLAND narrative-centered learning environment to incorporate adaptable event sequences, we conducted a pair of human subject studies to collect training data for inducing the narrative-centered tutorial planning models. It was necessary to conduct two studies, held at two different schools, to collect enough training data for inducing the models. We sought to adequately sample the space of possible narrative-centered tutorial planning policies, which required obtaining multiple observations for each combination of state and narrative adaptation. However, the data collection was also subject to the constraints of working with school systems, as well as resource constraints inherent in conducting multiple large (off-site) classroom studies. The studies were performed in classrooms within the participants’ schools. Prior to the study, the research team outfitted the

classroom with laptops and study materials in preparation for the data collection. The studies replaced students' regularly scheduled science classes on the days that students participated. While the study was part of students' classwork, the students did not receive a grade for their participation, and they were required to obtain parental consent before involvement. Students who did not receive permission to participate in the study completed alternate activities with one of their teachers. No personally identifiable information was collected. Each student was assigned a unique, anonymous identifier, which provided a mechanism for associating questionnaire and interaction data.

### **5.2.1 Population**

The first study involved 300 eighth-grade students from a rural North Carolina middle school. Among the students, 37 were removed from the dataset due to incomplete or inconsistent data. Data irregularities stemmed from absences and school-related interruptions during the multi-session study. In the remaining data, there were 130 males and 133 females. The students' ages ranged from 12 to 16 ( $M = 13.4$ ,  $SD = .58$ ). Students' ethnicities were approximately 69% Caucasian, 20% African American, and 10% Hispanic or Latino.

The second study involved 153 eighth-grade students from a different rural North Carolina middle school. Among the students, 14 were removed from the dataset due to incomplete or inconsistent data. There were 64 males and 75 females. The students' ages ranged from 13 to 15 ( $M = 13.6$ ,  $SD = .69$ ). Students' ethnicities were approximately 67% Caucasian, 21% Hispanic or Latino, and 11% African American. At the time of both studies, students had not yet completed the standard microbiology curriculum in their classes.

### **5.2.2 Design**

Every student across both studies used the same version of CRYSTAL ISLAND, which included the thirteen adaptable event sequences described in the previous section. Furthermore, students in both studies followed identical study procedures. There were no experimental conditions, and there was no control group. All students used CRYSTAL ISLAND individually. There were no inter-player interactions within the environment, as CRYSTAL ISLAND is

entirely single player. Students interacted with the environment until they solved the mystery, or 55 minutes of interaction time elapsed, whichever occurred first.

While using CRYSTAL ISLAND, students unknowingly encountered decision points several times. At each decision point, the environment selected a narrative adaptation according to a random policy, consistent with the modeling procedure presented in Chapter 3. By logging these narrative adaptations, as well as students' subsequent responses, the environment broadly sampled the space of policies for controlling adaptable event sequences.

In addition to adaptable event sequences, the CRYSTAL ISLAND environment included an in-game emotion self-report feature for a parallel affective computing study (Sabourin, Rowe, Mott, & Lester, 2011; Sabourin, Mott, & Lester, 2011). A subset of participants was also randomly selected to use CRYSTAL ISLAND while sitting on posture-sensitive seats. The posture-sensitive seats were designed to provide a minimally invasive measure of students' physiological state. Both of these features were designed to minimize their effect on the narrative-centered tutorial planning corpus collection, and they were not an active part of the current investigation. Therefore, they are not considered further in this work.

### **5.2.3 Pre-Study Materials**

Several questionnaires were administered prior to students using CRYSTAL ISLAND in order to gather data about individual characteristics and prior microbiology knowledge. Specifically, students completed questionnaires about demographics, video game-playing experience, microbiology content knowledge, personality, goal orientation, and emotion self-regulation strategies. The demographic survey included questions about students' gender, age, and ethnicity (Appendix A). The research team created the game-playing experience survey, which included several Likert items about students' perceived frequency and skill at playing video games (Appendix A). Both of these measures were the same ones used in the study of the non-adaptive CRYSTAL ISLAND described earlier. The microbiology content test was created through an iterative process by the research team, and eighth-grade science teachers reviewed the initial iterations for language and content. The test consists of 19 multiple-

choice items about viruses, bacteria, pathogens, mutagens, carcinogens, microscopes, infectious diseases, and the scientific method. The test was a refinement of the earlier microbiology curriculum test, and is shown in Appendix F.

The Big Five Inventory (BFI) was employed to measure students' personality characteristics in terms of five traits: openness, conscientiousness, extraversion, agreeableness, and neuroticism (John, Naumann, & Soto, 2008). The BFI is a valid, reliable measure consisting of 44-Likert items designed for efficient and flexible assessment of individual facets of personality. The Achievement Goals Questionnaire (AGQ) was employed to measure students' goal orientations, which characterize students' motivations in competence-relevant behavior (Elliot and McGregor, 2001). Achievement goals have traditionally been framed in terms of competence through task mastery and demonstrated competence relative to others. The AGQ is a 12-item measure for assessing a 2x2 framework comprised of mastery-approach, mastery-avoidance, performance-approach, and performance-avoidance achievement goal factors. The Cognitive Emotion Regulation Questionnaire (CERQ) was used to assess students' strategies for self-managing stressful emotions (Garnefski & Kraaij, 2006). The 18-item version of the CERQ was used, which includes two questions for each of nine subscales: Self-blame, Other-blame, Rumination, Catastrophizing, Putting into Perspective, Positive Refocusing, Positive Appraisal, Acceptance, and Planning.

Personality, goal orientation, and emotion regulation measures provide useful data about students' individual differences. Furthermore, these measures might also yield valuable features for driving narrative-centered tutorial planning decisions. However, they were primarily included in the two data collections to inform the parallel affective computing study. We leave investigation of these instruments for future work, and therefore we do not discuss them further in this dissertation.

#### **5.2.4 Post-Study Materials**

Students completed several questionnaires immediately following CRYSTAL ISLAND. The

questionnaires included measures of students' microbiology content knowledge, recall of the game's mystery narrative, intrinsic motivation, and presence. The microbiology content post-test consisted of the same 19 multiple-choice items that students completed prior to playing CRYSTAL ISLAND. The narrative questionnaire was a 13 item multiple-choice measure created by the research team (Appendix G). The questionnaire includes questions that assess participants' recall of key events and characters from the science mystery. The Intrinsic Motivation Inventory (IMI) was employed to measure students' intrinsic motivation, or subjective experience, toward CRYSTAL ISLAND (Ryan, 1982). Five sub-scales of the IMI were used: Interest/Enjoyment, Perceived Competence, Effort/Importance, Pressure/Tension, and Value/Usefulness. The IMI has been used in several experiments, and a study has found strong support for its validity (McAuley, Duncan, and Tammen, 1989). The Presence Questionnaire (PQ) was administered to assess students' self-reported presence while using CRYSTAL ISLAND (Witmer & Singer, 1998). Presence describes a user's sense of "being there" when interacting with a mediated environment, or alternatively the "the subjective experience of being in one place or environment, even when one is physically situated in another" (Witmer & Singer, 1998). Version 4.0 of the PQ was used, which consists of 33 Likert items across four subscales: Involvement, Sensory Fidelity, Adaptation/Immersion, and Interface Quality. Prior studies have found that self-reported presence is a significant predictor of student learning in CRYSTAL ISLAND (Rowe, Shores, Mott, & Lester, 2011).

### **5.2.5 Apparatus**

Students played CRYSTAL ISLAND on Lenovo Thinkpad notebook computers. Students used the WASD and/or arrow keys in order to control their avatar's movements in the virtual environment. Students used an external mouse and/or track pad to manipulate the gaze direction of the game's first-person camera. Students wore headphones in order to listen to the game's audio, including character speech, ambient island noises, and other sound effects. Valve Software's Source™ engine provided graphics rendering, audio playback, keyboard and mouse controls, and low-level character behavior functionalities.

### **5.2.6 Procedure**

The schools' science teachers notified students about the study. Students were given informed consent paperwork for their parents/guardians to sign, and the forms were returned to their science teachers several days later. Approximately one week prior to using CRYSTAL ISLAND, students completed the web-based pre-study questionnaires in a school computer lab. The questionnaires took approximately 30 minutes to complete.

The following week, students used CRYSTAL ISLAND during their science class period. For both studies, CRYSTAL ISLAND sessions were held in a classroom within the host school. Each student used CRYSTAL ISLAND for one session. Both studies were conducted over the course of several sessions, spanning multiple days at each school. The study classroom contained several dozen computers. The computers were positioned in order to discourage neighboring participants from viewing each other's screens. All computers hosted the same version of CRYSTAL ISLAND. Each station had several paper handouts to provide useful information about the study and virtual environment (Appendix E). The handouts included the following materials. A brief backstory was provided that informed students about the purpose of their arrival on CRYSTAL ISLAND, the characters and locations they would encounter on the island, and an explanation of some of the virtual equipment they would have access to in the story world. A cast of characters handout showed pictures of seven important characters that students would encounter in the game, as well as their names. A camp map handout provided an overhead screenshot of the island, with annotations marking key locations and buildings. A "controls" handout documented the keyboard and mouse keys that students would use to interact with virtual environment. The participants could reference the handout materials throughout the CRYSTAL ISLAND portion of the study.

After students were seated at their computers, a researcher provided a brief introductory presentation about the study's activities. The researcher also briefly described the handout materials and summarized the schedule for the study session. Students were given 55 minutes to use CRYSTAL ISLAND. Students used CRYSTAL ISLAND until they solved the mystery or the allotted time expired, whichever occurred first. Immediately after students

finished playing CRYSTAL ISLAND, they remained at their computers to complete the post-study questionnaires. The questionnaires were accessed through a web browser, and they took approximately 30 minutes to complete. After participants finished the post-study materials, they had a choice of “filler” activities: reading a book, completing a CRYSTAL ISLAND-themed word search, or sitting quietly. Students were instructed not to discuss the game or science mystery with their classmates until the study was over.

### **5.3 Structure of Narrative-Centered Tutorial Planning Corpus**

Using matched-pairs t-tests, an investigation of science learning gains found that students in both studies achieved significant improvements in microbiology content knowledge. In the first study, students improved from pre-test ( $M = 6.7$ ,  $SD = 2.4$ ) to post-test ( $M = 8.8$ ,  $SD = 3.4$ ),  $t(262) = 10.7$ ,  $p < .0001$ . Similarly, students in the second study improved from pre-test ( $M = 6.5$ ,  $SD = 2.2$ ) to post-test ( $M = 8.3$ ,  $SD = 3.4$ ),  $t(138) = 6.9$ ,  $p < .0001$ . In the first study, 34% of students solved CRYSTAL ISLAND’s science mystery in the time allotted. In the second study, 26% of students solved the mystery.

The data collected during both studies were combined into a single corpus for inducing narrative-centered tutorial planning policies. The corpus consisted of two parts: students’ interaction logs, and students’ pre/post questionnaire results. Although collected separately, the two types of data were mapped to one another using students’ unique identifiers. Only total scores on the questionnaires—not raw responses—were incorporated into the training corpus. Afterward, non-pertinent record types were removed from the interaction logs, including students’ position and orientation data, which were collected multiple times per second throughout the CRYSTAL ISLAND interactions.

The resulting data set consists of 315,407 observations, yielded by log data from 402 students. The corpus includes data on a range of student actions, including the following categories: conversing with virtual characters, reading books, viewing posters, picking up/dropping objects, smartphone usage, opening/closing doors, answering quiz questions, running laboratory tests, and filling out the diagnosis worksheet. The most commonly

observed student actions are student-agent dialog turns, opening doors, and viewing posters. Table 5 provides a more detailed breakdown of student action frequencies in the corpus.

In addition to student actions, there are 10,057 instances of narrative adaptations in the corpus, which correspond to approximately 25 narrative adaptations per student. Table 6 shows the break down of different narrative adaptations for each adaptable event sequence in order of frequency of occurrence<sup>7</sup>.

Table 5. Breakdown of student actions in the planning corpus.

Student Action	Total Count
Dialog Turn	68533
Open Door	24756
View Poster	20982
Respond to Quiz	19868
Pickup Object	17999
Drop Object	17986
Diagnosis Worksheet	14404
Use Smartphone	13838
Initiate Conversation	9594
Test Object	5408
Label Slide	3218
Take Note	1849
Read Book	1611
Use Bryce's Computer	1097
Close Door	886

<sup>7</sup> An error in CRYSTAL ISLAND's random number generator caused some uneven sampling of the Mystery Solution and Initial Lab Count AESs. However, the size of the overall data collection appears to have marginalized any impact the error may have posed for inducing the narrative-centered tutorial planner.

Table 6. Breakdown of narrative adaptations in the planning corpus.

Adaptable Event Sequence	Narrative Adaptation	Total Count	AES Proportion
Record Findings Reminder	Receive	1686	
	Not Receive	1749	
	Total	3435	34.16%
Next Goal Prompt	Receive	581	
	Not Receive	578	
	Total	1159	11.52%
Knowledge Quiz	Receive	547	
	Not Receive	526	
	Total	1073	10.67%
Diagnosis Worksheet Feedback	Minimal Detail	252	
	Moderate Detail	287	
	Maximum Detail	265	
	Total	804	7.99%
Details of Bryce's Symptoms	Minimal Detail	332	
	Moderate Detail	315	
	Total	647	6.43%
Details of Teresa's Symptoms	Minimal Detail	172	
	Moderate Detail	194	
	Maximum Detail	193	
	Total	559	5.56%
Quentin's Revelation	Deliver	222	
	Not Deliver	315	
	Total	537	5.34%
Bryce's Revelation	Deliver	212	
	Not Deliver	272	
	Total	484	4.81%
Reflection Prompt	Receive	224	
	Not Receive	231	
	Total	455	4.52%

Table 6 Continued

Adaptable Event Sequence	Narrative Adaptation	Total Count	AES Proportion
Mystery Solution	Influenza & Egg	51	
	Influenza & Milk	146	
	Influenza & Sandwich	49	
	Salmonellosis & Egg	41	
	Salmonellosis & Milk	63	
	Salmonellosis & Sandwich	60	
	<b>Total</b>		<b>410</b>
Initial Lab Test Count	Three Tests	206	
	Five Tests	85	
	Ten Tests	119	
	<b>Total</b>	<b>410</b>	<b>4.08%</b>
Increase Urgency	Deliver	192	
	Not Deliver	196	
	<b>Total</b>	<b>388</b>	<b>3.86%</b>
Discourage Off-Task Behavior	Receive	1	
	Not Receive	3	
	<b>Total</b>	<b>4</b>	<b>0.04%</b>
<b>TOTAL</b>		<b>10057</b>	

*Note:* The total counts for the 13 adaptable event sequences do not sum to the overall total, because the Bryce's Revelation AES and the Quentin's Revelation AES share decision points. Consequently, the "Not Deliver" narrative adaptations for these two AESs have been double counted.

This chapter has discussed the process of collecting a training corpus for inducing a narrative-centered tutorial planner for the CRYSTAL ISLAND learning environment. We described a modified version of CRYSTAL ISLAND's interactive narrative that incorporates thirteen AESs distributed across a range of narrative adaptation categories. Furthermore, we described a pair of studies involving middle school students using the modified version of CRYSTAL ISLAND with an exploratory narrative-centered tutorial planner. These studies were instrumental for generating a sizable corpus of student interactions, where students' problem

solving behaviors were shaped by randomly chosen narrative adaptations. The next chapter discusses how we utilize the corpus to model and solve a set of concurrent Markov decision processes that collectively compose the narrative-centered tutorial planner for CRYSTAL ISLAND.

## CHAPTER 6

### Evaluation Study

With the training corpus in hand, we investigate the efficacy of our framework for inducing narrative-centered tutorial planning models from student interaction data. A narrative-centered tutorial planner consisting of several concurrent narrative adaptation policies was devised to personalize adaptable event sequences in CRYSTAL ISLAND. In order to evaluate the impact of the induced narrative-centered tutorial planner on students' in-game problem-solving behavior, we conducted a controlled experiment with students from a North Carolina middle school. Students used two different versions of the CRYSTAL ISLAND environment, one version with the induced narrative-centered tutorial planner and another with a baseline planner. This chapter describes how the model induction procedure was implemented for CRYSTAL ISLAND, the evaluation study that was performed, and an empirical account of how the two versions of CRYSTAL ISLAND differed in their narrative adaptation behavior.

#### 6.1 An Implemented Narrative-Centered Tutorial Planner

In order to devise policies for tailoring narrative adaptations in CRYSTAL ISLAND, we modeled each of the narrative-centered learning environment's adaptable event sequences as independent Markov decision processes. In this case, "independence" supposes that actions performed under the guidance of one MDP policy will not affect the task environment of another MDP. This is a simplifying assumption, and one of its consequences is that all of the MDPs' task environments are considered static (i.e., non-changing) despite concurrent adaptations of other AESs. In addition to simplifying the model, the assumption obviates the need for collecting larger quantities of training data for modeling task environments that dynamically change in response to concurrent narrative adaptations. One approach for relaxing this assumption would be to induce narrative-centered tutorial planning policies

using on-line reinforcement learning techniques. However, in this work we focus on off-line methods for inducing planning policies. We leave an investigation of the conditions under which the independence assumption holds for future work.

Following the modeling procedure presented in Chapter 3, we identified a state representation, action sets, and a reward function for each AES in CRYSTAL ISLAND. Upon inspecting the training corpus we elected not to incorporate the Off-Task Discouragement AES. The training corpus included too few observations for inducing an Off-Task Discouragement narrative adaptation policy. This was likely a result of being too conservative in detecting off-task behaviors. Next, we formalized the 12 AESs as Markov decision processes. We used the training corpus described in Section 5.3 to calculate each MDP's environment model, and we utilized value iteration to optimize the planning policies. Additionally, we defined arbitration procedures for AESs in which concurrent policies could prescribe conflicting actions. Each of these steps is described in turn.

In the narrative-centered tutorial planning model, all of the concurrent MDPs shared the same state representation. While assigning distinct state representations to different AESs is a property of the framework presented in this dissertation, an extensive feature selection analysis was not possible due to scheduling limitations stemming from coordination with a public school system. Instead, we decided on a fixed set of state features, which included attributes empirically identified as useful in prior work, as well as features that characterized notable individual differences among students. Furthermore, it was necessary to identify features that could be calculated at any point in students' interactions.

The state representation consisted of eight binary features drawn from three categories: narrative features, individual difference features, and gameplay features. The state representation is shown in Figure 12. We limited the state representation to just eight binary features to reduce potential data sparsity challenges. The first four features were narrative-focused. Each feature was associated with a specific sub-goal from CRYSTAL ISLAND's narrative, and it indicated whether the sub-goal had been completed in the narrative thus far.

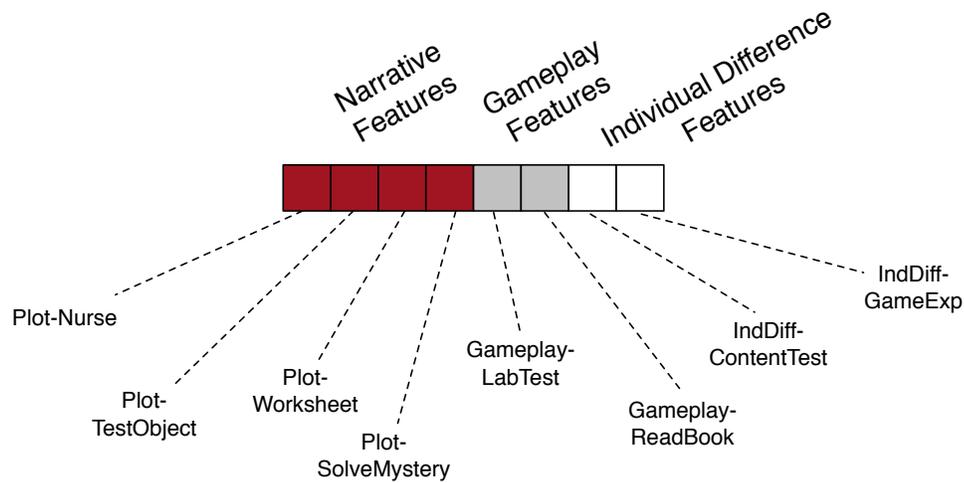


Figure 12. State representation for induced narrative-centered tutorial planner.

The four sub-goals were a subset of the eleven sub-goals described in Chapter 4. They included:

- **Camp nurse presents mystery’s main objective.** This sub-goal is accomplished when the student speaks to the camp nurse for the first time. The student is immediately asked to help identify the illness afflicting the research camp.
- **Test contaminated object in laboratory.** This sub-goal is accomplished when the student tests the transmission source for pathogenic contaminants using the laboratory’s testing equipment. The student is immediately informed that the food object is contaminated with some form of pathogen.
- **Submit diagnosis worksheet for review.** This sub-goal is accomplished when the student first attempts to submit her diagnosis worksheet to the camp nurse for review. Each worksheet submission is an attempt to solve the mystery. Submissions indicate that the student understands how to report her findings.

- **Submit correct, complete diagnosis.** This sub-goal is accomplished when the student solves the mystery by submitting a correct final diagnosis to the camp nurse.

These sub-goals were chosen because they distinguish between salient phases of the narrative problem-solving scenario, which emphasize different foci of gameplay actions. In theory, the four binary narrative features correspond to sixteen possible state configurations. However, only eight of the configurations can be achieved in CRYSTAL ISLAND due to ordering constraints among plot points.

The initial narrative state for each episode was [0, 0, 0, 0]. This state indicates that no sub-goals have been completed thus far. An intermediate state such as [1, 0, 1, 0] corresponds to an episode where the student has spoken to the camp nurse and he has attempted to submit his diagnosis worksheet for review. In this state, he has not yet tested the transmission source for pathogens or submitted a correct final diagnosis. There are two state configurations that are associated with a solved mystery. The configuration of [1, 0, 1, 1] describes a rare case where the student solved the mystery without testing the transmission source for pathogens. This solution is possible by guessing the solution based on patients' symptoms. Alternatively, the configuration of [1, 1, 1, 1] describes an episode where the student solved the mystery by completing all four sub-goals.

The next two state features were drawn from the individual difference category.

- **Microbiology Content Pre-Test Score.** This feature characterizes the student's prior knowledge of CRYSTAL ISLAND's science curriculum. It is calculated from the student's total score on the pre-study microbiology content test.
- **Gameplay Frequency.** This feature characterizes the student's prior experience playing video games, which is a proxy for the student's familiarity with the control scheme and first-person game genre. It is calculated from the self-reported gameplay frequency item in the pre-study demographic questionnaire.

These features were binary indicators calculated by performing median splits on students' pre-study microbiology content test scores and gameplay frequency self reports, respectively. Students who had at- or above-median values on these measures possessed values of 1 for the respective features. Students who had below-median values had values of 0 for the respective features. Both of the features' values were fixed throughout a given student's narrative-centered learning interaction. In other words, student actions in CRYSTAL ISLAND could not change the feature values. Incorporating these features required extending CRYSTAL ISLAND to support looking up pre-study questionnaire responses to calculate individual difference feature values for each student. The two individual difference features described above were chosen because microbiology content test performance has previously been shown to be a significant predictor of student learning and problem solving in CRYSTAL ISLAND, and gameplay experience varies considerably within (and between) student populations.

The final two state features were drawn from the gameplay category. Similar to the individual difference features, the gameplay features were also binary and required median splits to calculate their values. However, the values of these features could change during students' interactions with CRYSTAL ISLAND.

- **Laboratory Testing Behavior.** This feature characterizes whether the student is a frequent tester or an infrequent tester. The feature is calculated using a running count of the student's laboratory tests, which are compared to median testing frequencies for similar stages of the narrative-centered learning interaction. Specifically, we maintain a record of median test counts (calculated from the training corpus) at two-minute intervals. At each time interval, the student's current test count is compared to the median count for the same interval. If the student has tested objects with above-median frequency, the feature is assigned a value of 1. If the student has tested objects with at- or below-median frequency, the feature is given a value of 0.

- **Book Reading Behavior.** This feature characterizes whether the student is a frequent book reader, or an infrequent book reader. The feature is calculated using a running count of the student's books read, which are compared to median book reading frequencies for similar stages of the narrative-centered learning interaction. The details of its calculation are analogous to the laboratory testing behavior feature described above.

These gameplay features were chosen because they were previously identified as varying significantly across high-performing and low-performing groups in CRYSTAL ISLAND (Rowe, Shores, Mott, & Lester, 2010a). While the features were deemed promising, they represent a small subset of the range of gameplay features that could be calculated from students' gameplay characteristics. In particular, investigating alternate combinations of gameplay features for narrative-centered tutorial planning through feature selection methods is a promising direction for future work.

The action sets for the 12 MDPs corresponded to the narrative adaptations for the associated AESs. The narrative adaptations for each AES are shown in Chapter 5, Table 6. The cardinality of these actions sets ranged from binary decisions to 6-way decisions.

A single reward function was used for all of the MDPs, which was based on students' *normalized learning gains*. Prior work has demonstrated that learning, problem solving, and engagement measures are significantly correlated in CRYSTAL ISLAND (Rowe, Shores, Mott, & Lester, 2011). Furthermore, prior work has suggested that episodes concluding with the same final narrative state—a solved mystery—are often the episodes with greatest student learning and engagement. This observation motivated the decision to implement a single reward function across all MDPs. Normalized learning gains are calculated using students' pre- and post-study microbiology curriculum test scores. The formula for calculating normalized learning gains is shown in Equation 1. In order to determine the actual reward values, normalized learning gains were first calculated for each student, and then a median split was performed. Students who had normalized learning gains that were greater than or

equal to the median were awarded +100 points at the conclusions of their episodes. Students with normalized learning gains that were less than the median were awarded -100 points. The same reward design has yielded effective policies for alternate types of intelligent tutoring systems (Chi, VanLehn, & Littman, 2010). All rewards were associated with the final observed state transitions of each episode. In computing solutions to the MDPs, we sought to maximize the expected discounted reward obtained during an episode.

$$\text{Normalized Learning Gains} = \frac{\text{Posttest} - \text{Pretest}}{\text{Max Score} - \text{Pretest}} \quad (1)$$

The task environment for each MDP involved one or more decision points occurring each time a student satisfied the preconditions for an AES. The task environments were episodic, because students' interactive narrative experiences concluded when the mystery was solved or the time limit expired. The probability values in the state transition models  $P_i$  were calculated in terms of state values at successive decision points. In other words, state transitions did not characterize the immediate effects of narrative adaptations, but rather the aggregate impacts of narrative adaptations between sequential decision points for a given AES. If a particular AES in an episode possessed no successive decision point, the episode's final state was used instead.

The design of CRYSTAL ISLAND's narrative-centered tutorial planner posed one instance of simultaneous decision points for multiple AESs. Specifically, the Quentin's Revelation AES and the Bryce's Revelation AES could raise simultaneous decision points when triggered through a conversation with the camp nurse Kim. A conflict would arise between the two AESs if their corresponding policies both recommended delivering their respective revelations. Multiple plot revelations were not permitted within a single episode of CRYSTAL ISLAND's narrative. Consequently, it was necessary to implement an arbitration procedure for handling potential conflicts between the two AES policies. We elected to adopt a domain-independent procedure that favored the revelation with greatest Q-value (i.e., expected return) in the MDP model.

In order to obtain the narrative-centered tutorial planning policies, we used the *value iteration* algorithm (Sutton & Barto, 1998). The author implemented PyMDP, a reinforcement learning library developed in the Python programming language, which includes an implementation of the value iteration algorithm. The library provides reusable code for obtaining model-based solutions to MDPs, and it permits configurable state representations, action sets, and reward functions for use in a wide range of reinforcement learning applications. In addition to providing facilities for model-based solutions to MDPs, the library provides functionality for approximating state-transition and reward models from corpora of raw log data.

The 12 MDPs, one for each AES in CRYSTAL ISLAND, were implemented with the PyMDP library. The code base for the MDP investigation involved approximately 1600 lines of Python code. A discount factor of 0.9 was used in PyMDP's dynamic programming update rules. Using the PyMDP tool, each policy was learned in less than a minute. Distinct policies were induced for each AES. The resulting policies were encoded as direct mappings between state values and planner actions. The policies are shown in tabular form in Appendix H.

## 6.2 Evaluation Experiment

After inducing narrative-centered tutorial planning policies for each adaptable event sequence, we evaluated the planner's impact on student problem solving in the run-time CRYSTAL ISLAND environment. This required incorporating the induced narrative adaptation policies into CRYSTAL ISLAND by replacing the random narrative adaptation model from the corpus collection studies with the newly induced policies. The narrative adaptation policies were incorporated as external reference files, which were queried by CRYSTAL ISLAND whenever an adaptable event sequence was triggered. The narrative-centered tutorial planner resource files were human-readable and fully editable. The policies were not part of the CRYSTAL ISLAND binary, and they could be adjusted without re-compiling and re-linking CRYSTAL ISLAND's source code.

To evaluate CRYSTAL ISLAND's induced narrative-centered tutorial planner, we conducted a controlled experiment with middle school students, which compared the induced model to a baseline narrative adaptation approach. The sample of middle school students was drawn from a different school district than the one used for the corpus collection studies. The experiment involved taking out students from their regularly scheduled science classes in order to play CRYSTAL ISLAND.

### **6.2.1 Population**

A total of 75 eighth-grade students from a suburban North Carolina middle school participated in the study. Among these students, fourteen were removed from the dataset due to incomplete or inconsistent data. Data irregularities stemmed from absences and school-related interruptions during the multi-session study. Among the remaining 61 students, 31 were female and 30 were male. The students' ages ranged from 13 to 15 ( $M = 13.8$ ,  $SD = .62$ ). Students' ethnicities were approximately 38% Caucasian, 43% African-American, 15% Hispanic or Latino, 2% Alaskan or Native American, and 2% Asian. Students participated in the study as part of their everyday classwork, although their performance in the study did not impact their class grades. Students who did not receive permission to participate in the study completed alternate activities with one of their teachers.

### **6.2.2 Design**

The study had two conditions: a Baseline planner condition and an Induced planner condition. Students in both conditions played CRYSTAL ISLAND, but the conditions differed in terms of the narrative adaptation policies employed by each version's narrative-centered tutorial planner. The narrative-centered tutorial planners controlled events during the 12 adaptable event sequences in CRYSTAL ISLAND's interactive narrative. The Baseline Planner employed a uniform random policy to control how adaptable event sequences unfolded. In this condition, narrative adaptations were selected randomly whenever the planner encountered a decision point. This was equivalent to the exploratory narrative-centered tutorial planner used during the corpus collection studies, which were described in Chapter 5.

The Induced Planner followed policies obtained by solving the Markov decision processes associated with each adaptable event sequence, which were described in Section 6.1.

Students were randomly assigned to the two conditions when they entered the experiment room to play CRYSTAL ISLAND. Among students with complete data, random assignment resulted in 33 participants in the Induced Planner condition, and 28 participants in the Baseline Planner condition. All students played the game individually in the same large conference room. There were no inter-player interactions within the game. Students played until they solved the mystery or 70 minutes of interaction time expired, whichever occurred first.

### **6.2.3 Procedure**

School science teachers notified participants about the study. Participants were given informed consent paperwork for their parents/guardians to sign, and the forms were returned to their science teachers several days later. Approximately one week prior to the experiment, participants completed web-based pre-experiment questionnaires in a school computer lab. The questionnaires were identical to the measures used in the corpus collection study described in Section 5.2, and included the demographic questionnaire, gameplay experience questionnaire, microbiology content test, Big Five Inventory, Achievement Goals Questionnaire, and Cognitive Emotion Regulation Questionnaire. The pre-experiment questionnaires took approximately 30 minutes to complete.

The following week, participants left their regularly scheduled science classrooms in order to participate in the experiment. Each student participated in the experiment for only one session, but the entire experiment was conducted over several sessions during a single day. For each session, science teachers led the participants to the experiment room. Participants who had not returned an informed consent form moved to another classroom to complete alternate school activities with a different teacher. The experiment room contained several dozen computers. The computers were positioned several feet away from one another in order to discourage neighboring participants from viewing each others' screens. Half of

the computers hosted the Baseline Planner version of CRYSTAL ISLAND. The other half hosted the Induced Planner version of CRYSTAL ISLAND. Participants were randomly assigned to the two conditions by a researcher. After participants were seated at the computers, a researcher provided a brief introductory presentation about the experiment's activities. The researcher briefly described each of the handout materials, and summarized the schedule for the experiment session. Handouts were distributed that described the game's backstory, cast of characters, controls, and island map. The participants could reference these handout materials throughout the CRYSTAL ISLAND portion of the experiment.

Participants were given 70 minutes to play the game. Participants played CRYSTAL ISLAND until they solved the mystery or the allotted time expired, whichever occurred first. Immediately after participants finished playing CRYSTAL ISLAND, participants remained at their computers to complete the post-experiment questionnaires. The post-experiment questionnaires were identical to the ones used in the corpus collection studies, and they included the microbiology content test, narrative questionnaire, Intrinsic Motivation Inventory, and Presence Questionnaire. The questionnaires were accessed through a web browser, and they took approximately 30 minutes to complete. After participants finished the post-experiment materials, they had a choice of "filler" activities: reading a book, completing a CRYSTAL ISLAND-themed word search, or sitting quietly. Participants were instructed not to discuss the science mystery with their classmates until the experiment was over. Participants received no compensation, course credit, or grades for their involvement in the study.

### **6.3 Profile of Induced Narrative-Centered Tutorial Planner's Decisions**

The narrative adaptation policies in both versions of CRYSTAL ISLAND controlled events during the narrative-centered learning environment's 12 AESs. However, the policies do not provide a complete picture of how the two conditions varied in terms of interactive narrative behavior. During the experiment, students encountered narrative decision points with varying degrees of frequency, and they experienced narrative adaptations with a non-uniform distribution over planning states. Consequently, particular regions of the narrative-centered

tutorial planning policies were likely to be sampled more extensively than other regions. In order to provide a more complete understanding of how the two experimental conditions differed, we provide an empirical account of the narrative adaptations performed by the two planners in the evaluation experiment. The empirical account and the formal planning policies provide a detailed record of how the induced narrative-centered tutorial planner operated in the run-time CRYSTAL ISLAND environment. Hereafter, we refer to the induced narrative-centered tutorial planner as the Induced Planner, and the baseline exploratory narrative-centered tutorial planner as the Baseline Planner.

In aggregate, students in the Induced Planner condition experienced fewer narrative adaptations ( $M = 26.4$ ,  $SD = 6.6$ ) than students in the Baseline Planner condition ( $M = 29.6$ ,  $SD = 6.5$ ). A  $t$ -test indicates that the observed difference is marginally statistically significant,  $t(59) = 1.93$ ,  $p < .06$ . This difference is likely due to the Induced Planner's impact on student problem solving, which may have affected the number of narrative decision points that students encountered. In order to examine condition effects among the individual AESs, we consider the AESs in two groups: one-instance AESs and recurring AESs. The one-instance AESs (Quentin's Revelation, Bryce's Revelation, Initial Lab Test Count, Mystery Solution, and Increase Urgency) performed narrative adaptations just once during an episode. The recurring AESs (Diagnosis Worksheet Feedback, Record Findings Reminder, Knowledge Quiz, Next Goal Prompt, Details of Teresa's Symptoms, and Details of Bryce's Symptoms) could be encountered multiple times within an episode, and therefore could perform multiple narrative adaptations for a single student. We employ these groupings for methodological reasons, as distinct statistical methods are employed to examine the differences between conditions for each group.

We employ Pearson's chi-square test to examine differences among the five one-instance AESs across conditions. The two plot-revelation AESs (Quentin's Revelation and Bryce's Revelation) showed significant differences across conditions. The Induced Planner displayed a strong preference against delivering Quentin's Revelation, declining to deliver

Table 7. Contingency table for Quentin's Revelation AES.

	Quentin's Revelation		Total
	Not Deliver	Deliver	
Induced Planner	29 (88%)	4 (12%)	33
Baseline Planner	17 (61%)	11 (39%)	28
Total	46	15	61

Note.  $\chi^2(1, N = 61) = 6.15, p < .02$

Table 8. Contingency table for Bryce's Revelation AES.

	Bryce's Revelation		Total
	Not Deliver	Deliver	
Induced Planner	11 (33%)	22 (67%)	33
Baseline Planner	18 (64%)	10 (36%)	28
Total	29	32	61

Note.  $\chi^2(1, N = 61) = 5.91, p < .02$

the revelation in 88% of cases (Table 7). A Pearson's chi-square test indicates that the two experimental conditions differed significantly in their behavior toward the Quentin's Revelation AES. In contrast, the Induced Planner displayed a preference for delivering Bryce's Revelation, choosing the Deliver narrative adaptation in 67% of cases (Table 8). A chi-square test indicates that the Induced Planner's preference for delivering Bryce's Revelation was also significantly different from the Baseline Planner's behavior. In the contingency tables shown above for these AESs, interior cells contain occurrence counts for each narrative adaptation and experimental condition. Percentage values within parentheses indicate the relative proportions of each narrative adaptation for the given experimental

Table 9. Contingency table for Mystery Solution AES.

	Mystery Solution						Total
	Influenza & Egg	Influenza & Milk	Influenza & Sandwich	Salmonellosis & Egg	Salmonellosis & Milk	Salmonellosis & Sandwich	
Induced Planner	0 (0%)	0 (0%)	4 (12%)	29 (88%)	0 (0%)	0 (0%)	33
Baseline Planner	8 (29%)	2 (7%)	5 (18%)	7 (25%)	4 (14%)	2 (7%)	28
Total	8	2	9	36	4	2	61

Note.  $p < .001$

Table 10. Contingency table for Initial Lab Test Count AES.

	Initial Lab Test Count			Total
	3 Tests	5 Tests	10 Tests	
Induced Planner	4 (12%)	5 (15%)	24 (73%)	33
Baseline Planner	6 (22%)	12 (44%)	9 (33%)	27
Total	10	17	33	60

Note.  $p < .01$

condition (i.e., row).

In the Mystery Solution AES, the Induced Planner displayed an extremely strong preference for a single solution, choosing the Salmonellosis and Egg combination in 88% of student interactions (Table 9). In fact, the Induced Planner selected only one other solution—Influenza and Sandwich—during the entire experiment. A Fisher’s Exact Test indicates that the differences between the two experimental conditions in selecting Mystery Solution narrative adaptations are highly statistically significant.

In the Initial Lab Test Count AES, the Induced Planner displayed a strong preference for providing students the maximum number of initial tests (Table 10). The Induced Planner provided students the maximum number of initial lab tests during 73% of student interactions. A Fisher’s Exact Test indicates that the Induced Planner and Baseline Planner

performed narrative adaptations for the Initial Test Count AES significantly differently. We use Fisher's Exact Test rather than Pearson's chi-square test to account for cases of small cell frequencies.

In the Increase Urgency AES, the Induced Planner displayed a strong preference for the Not Delivering narrative adaptation. The urgency in virtual characters' dialog was increased in only 24% of cases in the Induced Planner condition, compared to 43% of cases in the Baseline Planner condition. However, a chi-square test does not find statistical evidence that the two experimental conditions were significantly different,  $\chi^2(1, N = 61) = 2.38, p = .12$ . Similarly, a chi-square test does not find statistical evidence of a condition effect for the Reflection Prompt AES,  $\chi^2(1, N = 61) = 1.69, p = .20$ .

In order to investigate the six recurring AESs, we employ two-tailed t-tests to compare the narrative adaptation frequencies for the two narrative-centered tutorial planners. We examine the total number of times that narrative adaptations were performed in each condition, as well as the relative proportions of narrative adaptations for each AES. The results of these comparisons are shown in Table 11. The values reported in the table are means per student, with standard deviation values shown in parentheses.

There were significant differences between the Induced Planner and Baseline Planner in the Diagnosis Worksheet Feedback AES. A t-test indicated that the two conditions differed significantly in the number of times they performed the Minimal Detail narrative adaptation,  $t(59) = 3.14, p < .01$ . Additionally, the two conditions demonstrated significant differences in their proportion of Minimal Detail narrative adaptations,  $t(50) = 2.93, p < .01$ , as well as proportion of Moderate Detail narrative adaptations,  $t(50) = -3.24, p < .01$ . In other words, the Induced Planner performed the Moderate Detail Diagnosis Worksheet Feedback narrative adaptation a larger fraction of the time than the Baseline Planner. Similarly, the Induced Planner performed the Minimal Detail narrative adaptation a smaller fraction of the time than the Baseline Planner. In general, the Induced Planner demonstrated a preference for moderately detailed feedback, choosing this narrative adaptation in 66% of Diagnosis Worksheet Feedback decision points.

Table 11. Breakdown of narrative adaptation frequencies for recurring AESs.

Adaptable Event Sequence	Narrative Adaptation	Total Count		Proportion of AES's Narrative Adaptations	
		Induced Planner	Baseline Planner	Induced Planner	Baseline Planner
Diagnosis Worksheet Feedback	Minimal Detail	<b>.18 (.58)</b>	<b>.82 (.98)</b>	<b>.07 (.21)</b>	<b>.27 (.27)</b>
	Moderate Detail	1.8 (1.8)	1.3 (1.7)	<b>.66 (.41)</b>	<b>.35 (.26)</b>
	Maximum Detail	.55 (1.0)	1.0 (.98)	.27 (.40)	.39 (.31)
	AES Total	2.5 (2.1)	3.1 (2.9)		
Record Findings Reminder	Receive	4.2 (3.6)	5.8 (3.0)	.44 (.28)	.48 (.15)
	Not Receive	5.5 (3.9)	6.2 (3.0)	.56 (.28)	.52 (.15)
	AES Total	<b>9.7 (4.3)</b>	<b>12.0 (4.8)</b>		
Knowledge Quiz	Receive	.85 (1.3)	1.3 (.90)	.33 (.48)	.46 (.31)
	Not Receive	1.8 (1.4)	1.4 (.84)	.67 (.48)	.54 (.31)
	AES Total	2.7 (.64)	2.7 (.71)		
Next Goal Prompt	Receive	<b>.67 (.99)</b>	<b>1.4 (1.1)</b>	<b>.23 (.35)</b>	<b>.46 (.33)</b>
	Not Receive	2.1 (1.4)	1.8 (1.3)	<b>.77 (.35)</b>	<b>.54 (.33)</b>
	AES Total	2.8 (1.3)	3.1 (1.3)		
Details of Teresa's Symptoms	Minimal Detail	.73 (.98)	.46 (.74)	.43 (.43)	.38 (.45)
	Moderate Detail	.48 (.71)	.32 (.55)	.30 (.41)	.32 (.45)
	Maximum Detail	.42 (.66)	.43 (.69)	.27 (.39)	.31 (.39)
	AES Total	1.6 (1.2)	1.2 (.96)		
Details of Bryce's Symptoms	Minimal Detail	<b>.24 (.56)</b>	<b>.64 (.87)</b>	<b>.17 (.34)</b>	<b>.41 (.44)</b>
	Moderate Detail	1.0 (1.1)	.96 (.88)	<b>.83 (.34)</b>	<b>.59 (.44)</b>
	AES Total	1.3 (1.3)	1.6 (1.1)		

Note: Statistically significant condition effects ( $\alpha = .05$ ) are emphasized in boldface.

Narrative adaptations for the Next Goal Prompt AES also differed significantly between the two planners. The Induced Planner delivered the prompt significantly fewer times than the Baseline planner,  $t(59) = 2.76, p < .01$ . Additionally, the Induced Planner delivered the prompt in a smaller proportion of decision points than the Baseline Planner,  $t(59) = 2.62, p < .02$ . In fact, the Induced Planner only delivered prompts about student goals

during 23% of decision points. This finding is surprising, as the Next Goal Prompt was expected to provide a useful scaffold for student problem solving. However, the Induced Planner used the narrative adaptation sparingly.

The Induced Planner aligned with expectations during AESs about Details of Bryce's Symptoms. The Induced Planner provided minimal detail about Bryce's symptoms significantly fewer times than the Baseline Planner,  $t(59) = 2.17, p < .04$ . Similarly, the Induced Planner offered minimal detail in a smaller proportion of decision points,  $t(50) = 2.18, p < .04$ , instead performing a larger proportion of narrative adaptations providing moderate detail on Bryce's symptoms. In fact, the Induced Planner elected to provide moderate detail during 83% of AESs about Details of Bryce's Symptoms.

No significant differences were observed for the AESs about Knowledge Quizzes or Details of Teresa's Symptoms. The Induced Planner did yield significantly fewer decision points about Record Findings Reminders,  $t(59) = 1.97, p = .05$ , but this was most likely an artifact of differences in student problem solving between the two conditions.

This chapter has described an implemented narrative-centered tutorial planner for the CRYSTAL ISLAND learning environment. We presented the state representation, action sets, and reward models for each of the constituent MDPs in the planner. Furthermore, we described an arbitration procedure used to resolve conflicts between narrative adaptation policies, as well as the methods used to induce planning policies from the training corpus of student interaction data. After describing the implemented planner, we presented an evaluation study with middle school students, which was designed to evaluate the induced narrative-centered tutorial planner's effectiveness at enhancing students' problem-solving processes compared to a baseline exploratory model. Finally, we provided an account of the differences in narrative adaptation behaviors between the two experimental conditions. Given these two narrative-centered tutorial planning models, the Induced Planner and the Baseline Planner, the next chapter presents results about their respective impacts on student learning and problem solving in CRYSTAL ISLAND.

## CHAPTER 7

### Results

In this chapter, we report empirical results on the effectiveness of the induced narrative-centered tutorial planner for enhancing student problem solving in CRYSTAL ISLAND. In order to evaluate the Induced Planner, we examine two categories of dependent measures: students' post-study questionnaire responses and students' in-game problem solving behaviors. In the latter case, we investigate condition effects on a range of gameplay metrics associated with problem-solving efficiency, as well as measures of deliberative problem-solving processes within CRYSTAL ISLAND. The findings reveal improved outcomes for students using the induced narrative-centered tutorial planner over a baseline approach, offering a demonstration of the effectiveness of our framework for modeling narrative-centered tutorial planning under uncertainty with concurrent Markov decision processes.

#### 7.1 Content Learning Results

Analyses of students' learning gains found students achieved significant improvements in microbiology content knowledge in both experimental conditions. In the Induced Planner condition, students significantly improved their content test scores by 1.6 questions on average from pre-test ( $M = 7.8$ ,  $SD = 2.2$ ) to post-test ( $M = 9.4$ ,  $SD = 3.6$ ),  $t(32) = 2.67$ ,  $p < .02$ . In the Baseline Planner condition, students also achieved significant improvements in content test score from pre-test ( $M = 7.2$ ,  $SD = 2.5$ ) to post-test ( $M = 9.5$ ,  $SD = 3.4$ ),  $t(27) = 4.09$ ,  $p < .001$ . A comparison between the two conditions' average post-test scores failed to find evidence of a significant condition effect on microbiology content learning. Similarly, no condition effects were observed on students' normalized learning gains, or individual items within the post-test.

## 7.2 Problem Solving Results

In order to investigate student problem solving in the two conditions, we first examined the rates at which students solved the mystery by successfully diagnosing the illness. In the Induced Planner condition, 61% of students solved the mystery, compared to 46% of students in the Baseline Planner condition. This finding is notable because students in the Induced Planner condition solved CRYSTAL ISLAND's mystery at a higher rate than in any prior study with comparable versions of the CRYSTAL ISLAND software<sup>8</sup>. Despite the sizable difference, a Pearson's chi-square test failed to find evidence that the difference between conditions was statistically significant. We hypothesize that this result was due to a lack of statistical power stemming from the study's small sample size. In order to follow up on the analysis, we examined condition effects on problem-solving efficiency. Specifically, we examined the time taken to solve CRYSTAL ISLAND's mystery in each experimental condition. Because several students in each condition did not complete the scenario in the allotted time, we devised a Projected Completion Time metric, which approximated completion times based on each student's number of remaining narrative sub-goals. In the calculation, we allocated five minutes for each remaining sub-goal. The calculation for the Projected Completion Time metric is shown in Equation 2.

$$\text{Proj. Completion Time} = \text{Session Time} + (300 \times \text{Unsolved Goal Count}) \quad (2)$$

On average, students in the Induced Planner condition had a Projected Completion Time of 3607 seconds ( $SD = 1286$ ), and students in the Baseline Planner condition had a Projected Completion Time of 4119 seconds ( $SD = 988$ ). A two-tailed t-test indicated that the difference between the two conditions was marginally statistically significant,  $t(59) = 1.72$ ,  $p = .09$ . These findings provide preliminary evidence that students interacting with the

---

<sup>8</sup> There were at least three large studies conducted with the modern version of CRYSTAL ISLAND prior to the evaluation experiment, including the two corpus collection studies. The mystery completion rates for those studies were 47%, 34%, and 26%, respectively. The first of the studies was conducted in the same school as the evaluation experiment, but it involved a different sample of students.

Induced Planner solved CRYSTAL ISLAND's story-centric problem-solving scenario more efficiently than students interacting with the Baseline Planner.

Next, we investigated several metrics that characterized students' problem-solving processes to examine whether students adopted different patterns of problem-solving behavior in the experiment's two conditions. In particular, we sought evidence of deliberate problem solving, which is contrasted with problem-solving strategies that involve extensive guessing or non-purposeful behavior. To perform this investigation, we calculated several metrics that have previously yielded insights in exploratory investigations of problem solving in CRYSTAL ISLAND, including measures of hypothesis testing efficiency (Spires, Rowe, Mott, & Lester, 2011) and early information gathering behavior (Sabourin, Rowe, Mott, & Lester, 2012).

We first analyzed students' hypothesis testing behaviors in the narrative-centered learning environment.<sup>9</sup> Students tested hypotheses about potential transmission sources by running tests in the camp's virtual laboratory. A two-tailed t-test indicated that students in the Induced Planner condition ( $M = 13.7$ ,  $SD = 10.9$ ) conducted marginally fewer tests than students in the Baseline Planner condition ( $M = 19.5$ ,  $SD = 14.4$ ),  $t(59) = 1.80$ ,  $p < .08$ . Additionally, the Induced Planner group ran significantly fewer tests ( $M = 4.7$ ,  $SD = 7.7$ ) after successfully identifying the transmission source than students in the Baseline Planner group ( $M = 11.0$ ,  $SD = 12.6$ ),  $t(59) = 2.39$ ,  $p < .03$ . These findings suggest that students in the Induced Planner condition were significantly more efficient in testing their hypotheses about the disease's transmission source.

Next, we examined student behaviors during the early stages of their investigations, i.e. prior to testing hypotheses about potential transmission sources in the virtual laboratory. A two-tailed t-test indicated that students in the Induced Planner condition completed significantly more narrative sub-goals prior to testing the disease's transmission source ( $M = 6.67$ ,  $SD = .85$ ) than students in the Baseline Planner condition ( $M = 5.82$ ,  $SD = 1.74$ ),

---

<sup>9</sup> We do not apply the Bonferonni correction because of the exploratory nature of our investigation of students' problem-solving processes.

$t(59) = -2.46, p < .02$ . In order to expand on this finding, we investigated how students collected background information on the scenario, as well as the microbiology curriculum, prior to forming, testing, and reporting hypotheses about the illness. We focus our analyses on three types of information gathering behaviors: conversations with virtual characters, examining virtual posters, and reading virtual books.

First, we investigated students' conversations with virtual characters. In terms of total number of conversations, as well as total number of dialogue turns, across the entire scenario, no significant differences between conditions were observed. Students in the Induced Planner condition initiated 19.8 conversations ( $SD = 8.4$ ) on average, and engaged in 79.2 total dialogue turns ( $SD = 27.7$ ). Students in the Baseline Planner condition initiated 22.0 conversations ( $SD = 8.1$ ) on average, and engaged in 86.9 total dialogue turns ( $SD = 28.3$ ). However, there were significant differences between conditions in student dialogue behavior with specific virtual characters. For example, students in the Induced Planner condition initiated significantly more conversations with Ford, the camp's virus specialist, prior to running a laboratory test,  $t(59) = -2.63, p < .02$ . Additionally, students in the Induced Planner condition engaged in significantly more dialogue turns with Ford prior to running a laboratory test,  $t(59) = -2.31, p < .03$ . Students in the Induced Planner condition also initiated marginally more conversations with Ford prior to submitting their diagnosis worksheet for review,  $t(59) = -1.71, p = .09$ . Conversely, Induced Planner students engaged in significantly fewer dialogue turns with Ford after running their first laboratory test,  $t(59) = 2.25, p < .03$ .

In combination, these findings suggest that students in the Induced Planner condition engaged in more extensive information gathering behaviors with Ford, the camp's virus specialist, prior to testing their hypotheses, rather than afterward. The trends in students' conversational behaviors with Ford are shown in Figure 13. In both of the figure's graphs, four phases of problem solving in CRYSTAL ISLAND are demarcated on the x-axis: 1) Actions prior to running the first lab test, 2) Actions prior to testing the disease's transmission source,

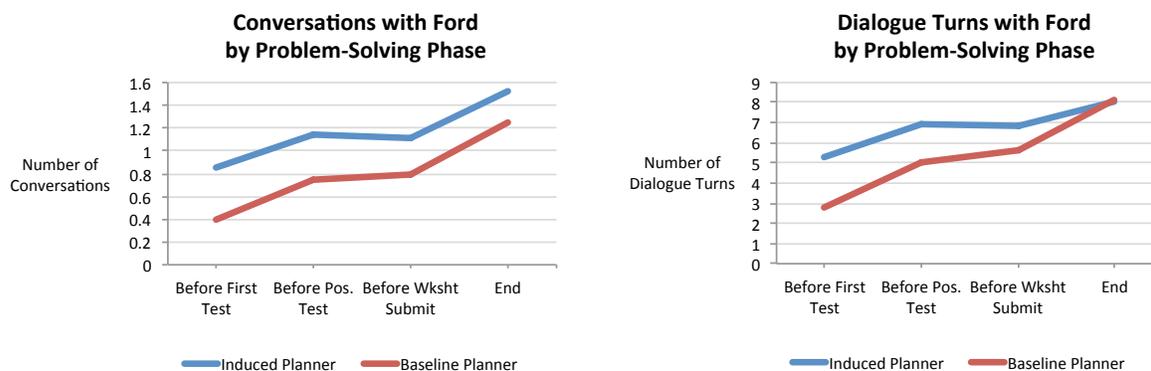


Figure 13. Students' conversational behaviors with the virus specialist.

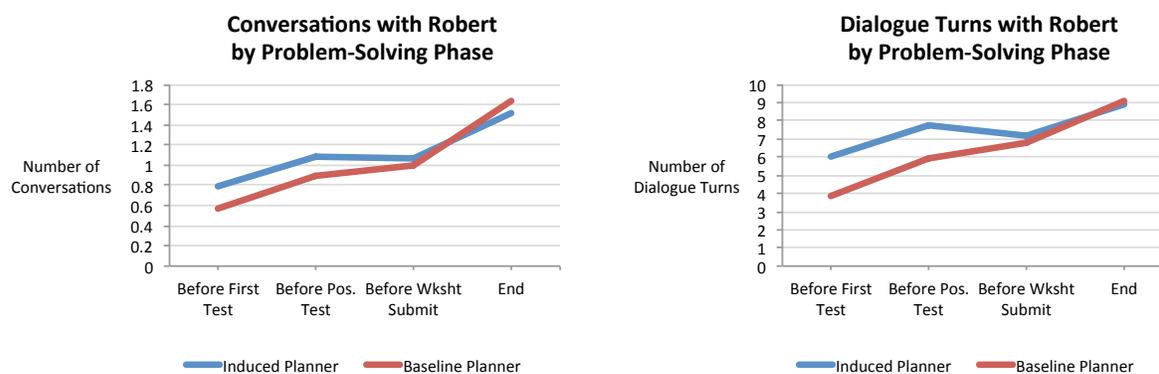


Figure 14. Students' conversational behaviors with the bacteria specialist.

3) Actions prior to first submitting the diagnosis worksheet for review, and 4) Actions prior to the scenario's conclusion. The y-axis indicates the total number of conversations, or dialogue turns, that the student performed during the associated problem-solving phase.

Similar patterns were observed in students' conversational behaviors with Robert, the camp's bacteria specialist (Figure 14). Students in the Induced Planner condition engaged in marginally more dialogue turns with Robert prior to running a laboratory test,  $t(59) = -1.71$ ,  $p = .09$ . Students in the Induced Planner condition also engaged in significantly fewer dialogue

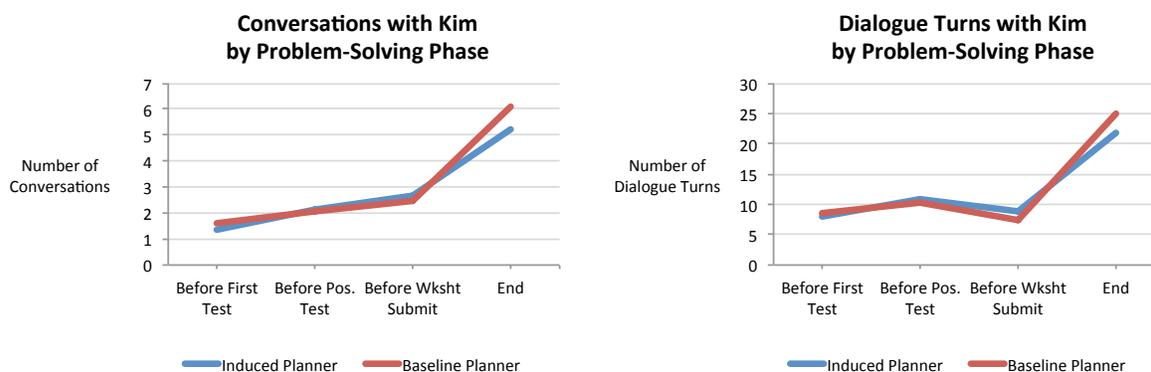


Figure 15. Students' conversational behaviors with the camp nurse.

turns with Robert after running their first laboratory test,  $t(59) = 2.21, p < .04$ . In a related finding, the Induced Planner students engaged in significantly more dialogue turns with the camp nurse Kim—a character who provides general background on pathogens, mutagens, and carcinogens—prior to first submitting their diagnosis worksheet,  $t(59) = -2.19, p < .04$  (Figure 15). These patterns in dialogue behavior are consistent with deliberate problem solving in CRYSTAL ISLAND. The findings suggest that students in the Induced Planner condition collected more background information about microbiology concepts prior to testing hypotheses about the illness, and students in the Baseline Planner condition gathered that background information afterward.

We also found differences between conditions in information gathering behavior related to the problem-solving scenario (as opposed to microbiology content). First, we investigated dialogues with the Bryce character, the sick lead scientist who provides information about his symptoms (Figure 16). Students in the Induced Planner condition initiated significantly fewer conversations with Bryce after running their first laboratory test,  $t(59) = 2.36, p < .03$ . Additionally, the Induced Planner group engaged in significantly fewer dialogue turns with Bryce after running their first laboratory test,  $t(59) = 2.08, p < .05$ . In examining students' conversational behavior with Quentin, the camp cook who provides

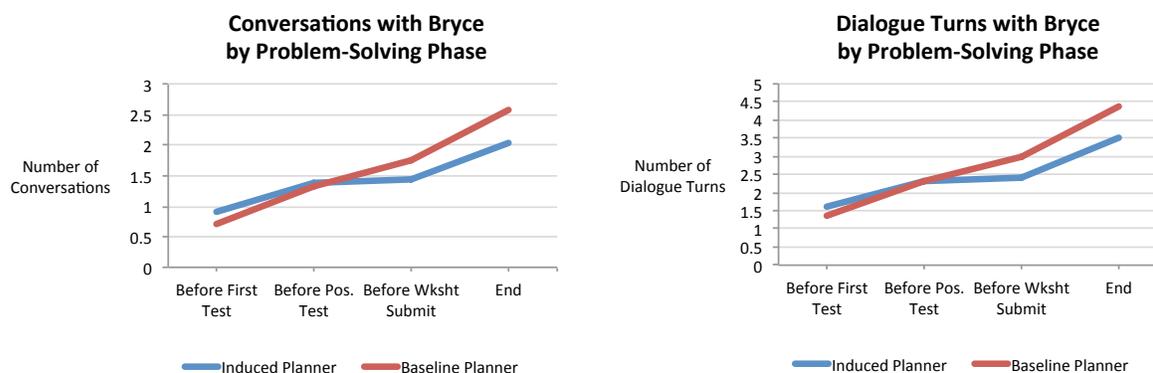


Figure 16. Students conversational behaviors with the sick lead scientist.

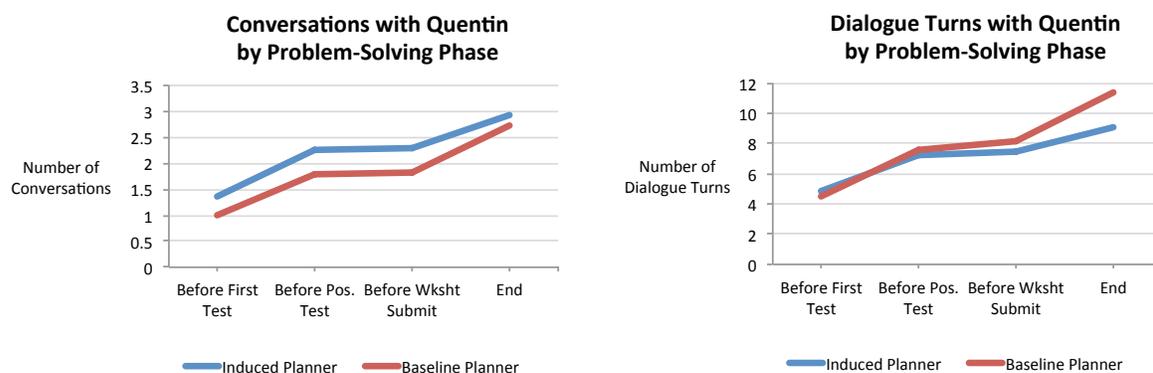


Figure 17. Students' conversational behaviors with the camp cook.

information about the team's recent eating habits, a chi-square test indicated that students in the Induced Planner condition were significantly more likely to converse with Quentin before running their first laboratory test than students in the Baseline Planner condition,  $\chi^2(1, N = 61) = 4.37, p < .04$  (Figure 17). However, students in the Induced Planner condition later engaged in marginally significantly fewer dialogue turns with the camp cook after running their first laboratory test,  $t(59) = 1.77, p < .09$ . There were no significant differences observed in conversational behaviors with the remaining virtual characters, the sick scientist Teresa and the laboratory technician Elise.

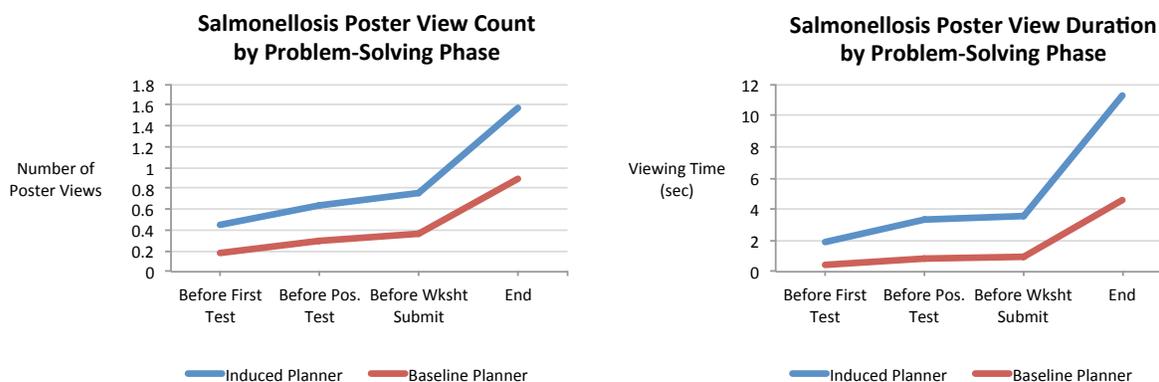


Figure 18. Students' Salmonellosis poster-viewing behaviors.

As a further investigation of students' information gathering strategies, we examined poster-viewing behaviors between experimental conditions. In these analyses, we only considered instances lasting longer than one second in duration, in order to avoid tallying inadvertent glances at the posters distributed throughout the virtual environment. Similar to the character dialogue findings, no significant differences in total poster viewing metrics were observed between conditions. The Induced Planner group viewed posters 16.6 times ( $SD = 16.0$ ) on average, for a total duration of 123.4 seconds ( $SD = 122.9$ ). Students in the Baseline Planner condition viewed posters 14.8 times ( $SD = 9.5$ ) on average, for a total duration of 122.4 seconds ( $SD = 118.9$ ). However, in an examination of the camp's six disease-focused posters, two-tailed t-tests indicated that students in the Induced Planner condition spent significantly more time reading the Salmonellosis poster prior to submitting their diagnosis worksheet than students in the Baseline Planner condition,  $t(59) = -2.18, p < .04$  (Figure 18). Similarly, students in the Induced Planner condition viewed the Anthrax poster more times prior to submitting their diagnosis worksheet, an effect that was marginally statistically significant,  $t(59) = -1.67, p = .1$  (Figure 19). Students in the Induced Planner condition viewed the Botulism poster more times prior to successfully testing the transmission source in the laboratory, an effect that was marginally statistically significant,

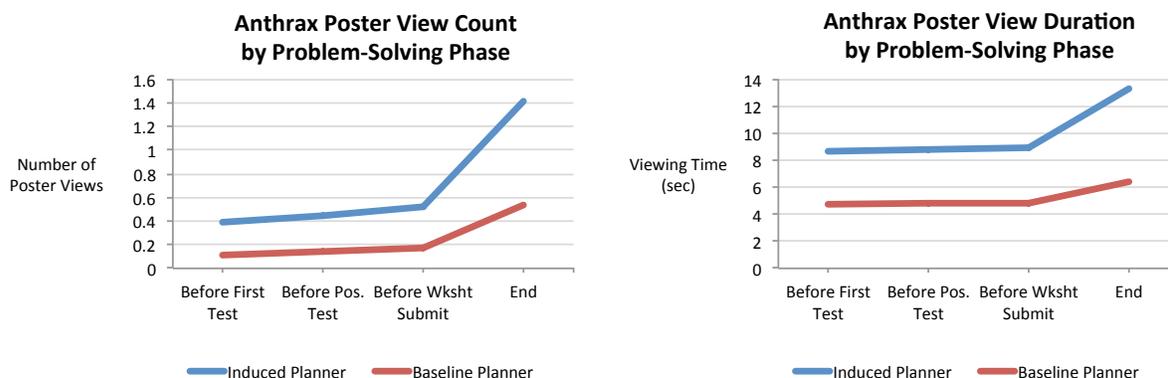


Figure 19. Students' Anthrax poster-viewing behaviors.

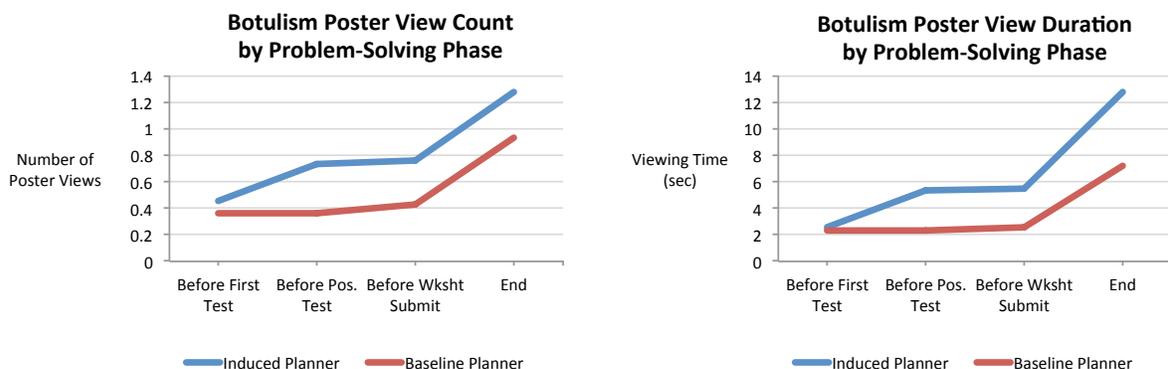


Figure 20. Students' Botulism poster-viewing behaviors.

$t(59) = -1.73, p < .09$  (Figure 20). Students in the Induced Planner condition also viewed the Ebola poster significantly more times prior to submitting their diagnosis worksheet,  $t(59) = -1.96, p = .05$  (Figure 21). No significant condition effects were observed for the Influenza or Smallpox posters.

Students in the Induced Planner condition viewed the viral disease poster marginally significantly more times prior to submitting their diagnosis worksheet than the Baseline Planner condition,  $t(59) = -1.93, p < .06$ . Conversely, students in the Baseline Planner



Figure 21. Students' Ebola poster-viewing behaviors.

condition viewed the viral disease poster more times after testing the transmission source in the laboratory, an effect that was marginally statistically significant,  $t(59) = 1.98, p < .06$ . Students in the Baseline Planner condition viewed the scientific method poster more times after first submitting their diagnosis worksheet, an effect that was marginally significant,  $t(59) = 1.84, p < .08$ . The Baseline Planner group also viewed the scientific method poster for marginally longer durations after first submitting their diagnosis worksheet,  $t(59) = 1.78, p < .09$ . The Baseline Planner group viewed the Are Viruses Alive poster more often than the Induced Planner group prior to testing the transmission source in the laboratory, an effect that was marginally significant,  $t(59) = 1.95, p < .06$ . Students in the Baseline Planner condition viewed the bacteria structure poster for longer durations than the Induced Planner group after running their first test in the laboratory, an effect that was also marginally statistically significant,  $t(59) = 1.83, p < .08$ . Interestingly, a similar pattern was not observed for the virus structure poster.

In aggregate, these findings suggest that students in the Induced Planner condition examined disease-specific posters more frequently before testing hypothesized diagnoses, particularly in the cases of posters about bacterial diseases. However, there were several cases in which Baseline Planner students examined posters about general microbiology concepts more frequently, and for greater durations, later in their investigations.

The findings raise questions about whether similar patterns were observed for students reading virtual books, which provide similar information for diagnosing the illness. However, an investigation of virtual book-reading behaviors did not resemble students' poster-viewing habits. We did not find compelling evidence of significant differences between conditions for book reading frequencies, or book reading durations, including among specific books and problem-solving phases. Students in the Induced Planner condition opened 5.7 books ( $SD = 6.0$ ) on average, and they read virtual books for 121.3 seconds ( $SD = 167.4$ ). Students in the Baseline Planner condition opened 5.4 books ( $SD = 4.8$ ) on average, and they read virtual books for 170.0 seconds ( $SD = 245.3$ ).

Significant condition effects were not observed for students' diagnosis worksheet behaviors. The diagnosis worksheet provided a resource for students to test their hypothesized diagnoses, but there was little evidence to suggest that the two conditions differed in terms of diagnosis worksheet efficiency or accuracy. We did not find evidence of significant condition effects on the number of fields students changed on their diagnosis worksheets,  $t(59) = -.74, p = .46$  (Induced Planner:  $M = 34.5, SD = 16.6$ ; Baseline Planner:  $M = 31.2, SD = 18.1$ ), nor the number of times they accessed their diagnosis worksheets,  $t(59) = -.58, p = .57$  (Induced Planner:  $M = 11.3, SD = 7.9$ ; Baseline Planner:  $M = 10.2, SD = 6.7$ ). Across the entire mystery scenario, the Induced Planner group submitted their diagnosis worksheets for review 2.5 times ( $SD = 2.1$ ) on average, and the Baseline Planner group submitted their diagnosis worksheets 3.1 times ( $SD = 2.9$ ), a difference that was not observed to be statistically significant,  $t(59) = .87, p = .39$ . Using a scoring rubric devised by Shores, Rowe, & Lester (2011), we examined the accuracies of students' final worksheets. Students in the Induced Planner condition achieved average final worksheet scores of 84.6 ( $SD = 33.7$ ), and students in the Baseline Planner condition achieved average final worksheet scores of 75.7 ( $SD = 45.0$ ). The difference in accuracies between the two conditions was not observed to be statistically significant,  $t(59) = -.89, p = .38$ . Further analyses of students' diagnosis worksheet behaviors by problem-solving phase and worksheet section also failed to reveal significant condition effects.

### 7.3 Engagement Results

In addition to investigating students' learning and problem-solving outcomes, we examined students' engagement-related questionnaire data from the post-study measures. An investigation of students' responses to the Presence Questionnaire did not find evidence of a significant difference between conditions,  $t(57) = -.08$ ,  $p = .93$ . Students in the Induced Planner condition reported an average PQ score of 143.2 ( $SD = 27.1$ ), and students in the Baseline Planner condition reported an average PQ score of 142.6 ( $SD = 23.9$ ).

Similar results were observed in students' responses on the Intrinsic Motivation Inventory. On the Interest/Enjoyment sub-scale, students in the Induced Planner group averaged a score of 4.5 ( $SD = 1.6$ ), as did students in the Baseline Planner group ( $M = 4.5$ ,  $SD = 1.3$ ). Unsurprisingly, the two groups were not observed to be significantly different,  $t(59) = .08$ ,  $p = .93$ . On the Perceived Competence sub-scale, students in the Induced Planner group averaged a score of 4.9 ( $SD = 1.3$ ), and students in the Baseline Planner group averaged a score of 4.7 ( $SD = 1.4$ ), a difference that was also not observed to be statistically significant,  $t(59) = -.40$ ,  $p = .69$ . There were no observed significant differences on the Effort/Importance sub-scale,  $t(59) = .06$ ,  $p = .95$ , where students in the Induced Planner group averaged 4.7 ( $SD = 1.4$ ) and students in the Baseline Planner group averaged 4.8 ( $SD = 1.5$ ). Similarly, there were no observed differences between the Induced Planner group ( $M = 2.4$ ,  $SD = 1.3$ ) and Baseline Planner group ( $M = 2.8$ ,  $SD = 1.2$ ) on the Pressure/Tension sub-scale,  $t(59) = 1.09$ ,  $p = .28$ . Finally, no significant differences between conditions were observed for the Value/Usefulness sub-scale,  $t(59) = -.09$ ,  $p = .93$  (Induced Planner:  $M = 5.3$ ,  $SD = 1.4$ ; Baseline Planner:  $M = 5.2$ ,  $SD = 1.7$ ).

We also examined students' recall of CRYSTAL ISLAND's narrative elements, such as key events and character roles. Students in the Induced Planner condition averaged a score of 10.5 ( $SD = 2.3$ ) on the narrative recall measure, and students in the Baseline Planner condition averaged a score of 9.9 ( $SD = 2.5$ ). The difference between the two groups did not achieve statistical significance,  $t(59) = -1.13$ ,  $p = .26$ . However, further item-level analyses found evidence of a condition effect on one of the narrative measure's items: question 10.

This question asked what action was most appropriate to take in CRYSTAL ISLAND's story after identifying the transmission source in the laboratory. The correct answer was, "Ask the lab technician to view the object under a microscope." While 79% of students in the Induced Planner group answered this question correctly, only 46% of students in the Baseline Planner group did so. The difference between the two conditions was statistically significant,  $\chi^2(1, N = 61) = 6.98, p < .01$ . This result mirrors the difference in testing behaviors between the two conditions after identifying the disease's transmission source.

This chapter has presented empirical results from an experiment to evaluate the induced narrative-centered tutorial planner's effectiveness at enhancing students' problem-solving processes in CRYSTAL ISLAND. We found that students interacting with the Induced Planner demonstrated greater efficiency in completing CRYSTAL ISLAND's problem-solving scenario than students in the baseline condition, as indicated by scenario completion times. Furthermore, we found that students in the Induced Planner condition exhibited improved patterns of deliberate problem-solving behavior. Specifically, students who interacted with the Induced Planner completed more narrative sub-goals prior to the hypothesis-testing phase of solving the mystery. Furthermore, they performed more information gathering behaviors (e.g, conversations with virtual characters, viewing disease-focused posters) prior to testing hypotheses about the disease and transmission source, as opposed to the Baseline Planner group that performed more of these behaviors after hypothesis testing. Finally, students in the Induced Planner group demonstrated greater efficiency during hypothesis testing, as they performed fewer laboratory tests in total, and fewer tests after identifying the transmission source, than students in the Baseline Planner condition. These patterns of deliberate problem solving are contrasted with undesirable problem-solving strategies involving extensive guessing or non-purposeful behavior.

These results provide substantial evidence that the Induced Planner successfully enhanced students' problem-solving processes in CRYSTAL ISLAND. However, we did not find evidence that the Induced Planner also enhanced students' learning and engagement outcomes, as initially hypothesized. In the next chapter, we propose explanations for how the

Induced Planner's specific narrative adaptations impacted students' problem-solving behaviors and learning outcomes. Furthermore, we discuss specific implementation decisions made in modeling the narrative-centered tutorial planner that may have impacted the findings.

## CHAPTER 8

### Discussion

Enhancing students' problem-solving processes in narrative-centered learning environments is a key challenge for narrative-centered tutorial planners. Results from the evaluation experiment indicate that our framework for inducing narrative-centered tutorial planners with concurrent Markov decision processes produces models that improve students' problem-solving efficiency, as well as patterns of deliberate problem-solving, in a narrative-centered learning environment. Specifically, fine-grained analyses of students' gameplay behaviors—such as conversing with virtual characters, viewing virtual posters, and running virtual laboratory tests—have revealed significant impacts of the narrative-centered tutorial planning policies induced from prior students' interaction data relative to a baseline exploratory planner. Because solving CRYSTAL ISLAND's science mystery demands a range of self-regulated learning skills, such as goal setting and strategy use, the results suggest that decomposing narrative-centered tutorial planning tasks into AESs is a promising framework for producing tailored story-based learning experiences that are both effective and engaging. It is also particularly encouraging that narrative-centered tutorial planning models induced from student interactions in a rural school district were effective at impacting a substantially different population of students: a highly diverse suburban school district.

In this chapter, we interpret results from the evaluation experiment and describe implications of the work. We propose explanations for the specific mechanisms that shaped students' problem-solving processes in the two experimental conditions. In addition, we explain why we likely did not find evidence of significant condition effects on students' microbiology content test scores or engagement-related measures. We also discuss limitations of the work, both in terms of the evaluation methods employed, as well as

limitations in the narrative-centered tutorial planning framework that is the focus of this dissertation.

### **8.1 Enhancing Students' Problem-Solving Processes with Narrative Adaptations**

While the evaluation experiment's results demonstrate the promise of our framework, they also raise questions about which specific mechanisms in the induced narrative-centered tutorial planner impacted students' problem-solving processes. In order to begin answering these questions, it is instructive to consider the aggregate differences in narrative adaptation behaviors between the two experimental conditions. As described in Chapter 6, the Induced Planner displayed a significantly greater inclination to deliver Bryce's revelation, provide the maximum number of initial laboratory tests, select the Salmonellosis & Egg mystery configuration, provide moderately detailed diagnosis worksheet feedback, and offer moderately detailed descriptions of Bryce's symptoms. In contrast, the Baseline Planner displayed a significantly greater inclination to deliver Quentin's revelation (although still not high at just 39% of opportunities), provide minimal diagnosis worksheet feedback, prompt the student toward the next goal, and provide minimal details about Bryce's symptoms. Many of the problem-solving metrics that differed between the two conditions related to students' information gathering behaviors prior to hypothesis testing, as well as subsequent hypothesis testing efficiency. Among the narrative adaptations preferred by the induced narrative-centered tutorial planner, only a subset could have plausibly produced these effects: the Bryce's Revelation AES, the Initial Test Count AES, and the Bryce's Symptoms AES. These are the only AESs that could both occur during early stages of the mystery as well as actively shape students' early problem-solving behavior. Among these three candidates, the Bryce's Revelation AES seems most likely. This is because the Initial Test Count AES did not occur until the student performed her first laboratory test (and hence it could not have affected prior gameplay behaviors), and the Bryce's Symptoms AES simply detailed an extra symptom of the disease, information that could have also been obtained from other virtual characters.

A majority of students in the Induced Planner condition received Bryce's revelation. However, if Bryce's revelation directly affected students' problem-solving processes as we observed in the prior chapter, students would need to encounter the revelation early in their investigations, prior to their main hypothesis testing phase of solving the mystery. Furthermore, there would need to be a plausible explanation about why Quentin's revelation did not have a similar effect for the Baseline condition. Students in the Baseline Planner condition were more likely to receive Quentin's revelation in aggregate, which served a similar tutorial purpose to Bryce's revelation: providing a low-to-moderate level hint about the mystery's solution.

The narrative adaptation statistics presented in this dissertation thus far have only measured the planners' aggregate behaviors; they have not provided insights about *when* students encountered particular narrative adaptations. However, we can address this gap by examining student encounters with narrative adaptations during their early phases of problem solving. A follow-up examination of the Bryce's Revelation AES and Quentin's Revelation AES indicated that students in the Induced Planner group did in fact encounter Bryce's revelation prior to hypothesis testing more often than the Baseline Planner group,  $\chi^2(1, N = 61) = 4.49, p < .04$ . Further, we failed to find evidence of a significant difference between conditions in whether students encountered Quentin's revelation before hypothesis testing,  $\chi^2(1, N = 61) = 1.85, p = .17$ . These findings are consistent with an explanation that the Bryce's Revelation AES had a key role in the observed impacts on students' early problem-solving behaviors between conditions, particularly those behaviors preceding initial tests in the laboratory. Furthermore, they are consistent with the premise that Bryce's revelation may have shaped the observed differences in hypothesis testing efficiency.

However, we must still consider *how* Bryce's revelation may have impacted student problem solving. As previously described, the Bryce's Revelation AES informed students about a nearby computer in the virtual environment that contained an e-mail message with details about Bryce's earlier investigation of the disease. While the e-mail included revelatory clues about the mystery's solution, our data indicate that when students viewed the

e-mail message they did not fully internalize the findings in such a way to immediately go solve the mystery. Instead, the hint may have served more as a scaffold, giving students an initial mental framework for pursuing and interpreting relevant information that they encountered in the virtual environment. Alternatively, the hint may have served as a priming mechanism, raising students' awareness of key facets of the mystery's solution, even though they did not fully grasp the message's content.

In support of the scaffolding explanation, consider that Bryce's revelation encouraged students to primarily focus on examining disease-specific background information, because the e-mail directly revealed which food object was the transmission source and whether the infection was bacterial or viral in nature. This content might explain why students in the Induced Planner condition demonstrated greater frequency, and duration of time, in reading disease-focused posters prior to hypothesis testing. A similar mechanism may have also produced the differences in observed conversational behavior with the camp's virtual microbiology specialists.

However, the priming explanation might account for differences in students' hypothesis testing behaviors. Students in the Induced Planner condition ran significantly fewer tests after testing the transmission source than students in the Baseline Planner condition. Perhaps the Bryce's Revelation AES made students in the Induced Planner condition more attuned to positive test results related to the transmission source, at which point they chose to proceed to the next stage of problem solving. In contrast, students in the Baseline Planner condition would have lacked this same awareness, continuing to run laboratory tests after successfully testing the disease's transmission source object. Speculating further, students in the Baseline Planner condition may have struggled to effectively interpret their own findings, leading to gathering information that was not directly essential for solving the mystery. This would account for findings that the Baseline Planner group demonstrated increased frequency and time viewing the Are Viruses Alive poster and Bacteria Structure poster at various stages of the mystery. Further, they may have sought out strategy-level information for clues on how to systematically diagnose the illness, such as the

Scientific Method poster. This additional attention allocated toward curricular content, rather than information immediately critical to diagnosing the illness, may also account for the small, but not statistically significant, elevation in learning gains observed for the Baseline Planner condition.

We were initially surprised that we did not find evidence that the Induced Planner was more effective than the Baseline Planner at enhancing students' recall of microbiology content after using CRYSTAL ISLAND. However, upon closer examination of the two planners, as well as students' gameplay behaviors, the likely reason for this result becomes clearer. The differences in the two planners' narrative adaptation behavior focused on shaping how students' solved the mystery scenario, as opposed to how they cognitively processed microbiology content. When originally designing the AESs for CRYSTAL ISLAND, the following AESs were devised to enhance students' content learning, either directly or indirectly: Record Findings Reminder, Knowledge Quiz, Next Goal Prompt, and Initial Lab Test Count. These AESs were intended to reduce students' cognitive load, encourage active processing of microbiology concepts, increase exposure to microbiology-focused resources in the environment, and promote effective application of microbiology knowledge. However, the policies for controlling most of these AESs either were not observed to differ significantly between the two planners, or the Induced Planner displayed a preference for *not* providing the associated coaching. For example, the Induced Planner delivered the Next Goal Prompt during just 23% of opportunities, and presented Knowledge Quizzes during just 33% of opportunities. Further, by providing students with a higher number of initial laboratory tests in the Induced Planner condition, students likely did not need to answer as many microbiology quiz questions to earn additional tests for identifying the disease's transmission source. These quizzes were intended to provide a mechanism for students to cognitively process microbiology content, as well as self-monitor their understanding of the curriculum. The Induced Planner's inclination to provide the maximum number of initial tests may have had a side effect of promoting reduced cognitive processing of microbiology content, with the tradeoff of greater efficiency in solving the mystery.

A notable trend in the experiment's results was that students in the Baseline Planner condition investigated several microbiology-focused posters with greater frequency and duration than students in the Induced Planner condition. A possible explanation for this result is students needed to improve their microbiology content knowledge to correctly answer quiz questions to earn additional laboratory tests. Because students in the Baseline Planner condition were allocated fewer initial lab tests, they had an added incentive to review microbiology concepts to improve their ability to earn additional lab tests. In contrast, students in the Induced Planner condition could focus more on identifying the transmission source and solving the mystery, rather than answering quiz questions about microbiology content.

This speculation raises questions about why the Induced Planner did not adopt policies that resulted in improved microbiology content knowledge compared to the Baseline Planner. After all, the induced narrative-centered tutorial planner sought to optimize normalized learning gains during training. A plausible explanation is that the particular state representation we employed for inducing narrative-centered tutorial planning policies may have had a significant effect on this form of behavior. In order to mitigate data sparsity issues, we selected narrative state features that provided "snapshots" of CRYSTAL ISLAND's plot. In contrast, we could have employed a narrative state representation that provided a history of plot points to-date, including details about event ordering. We chose the former state representation because it substantially reduced the number of state configurations considered by the planning model. However, the same decision also created opportunities for the narrative-centered tutorial planning task to deviate from the framework's Markovian independence assumptions. In the training corpus, the "solved mystery" state was associated with superior content learning outcomes, and was likely associated with learning progressions that similarly maximized students' cognitive processing of microbiology material. However, the state representation used in the Induced Planner did not directly model these narrative progressions, and thus it could induce policies that would efficiently

move students toward the solved mystery state without optimizing for paths that optimize content retention.

In principle, it would be straight-forward to modify the state to incorporate a narrative representation that encapsulates the history of plot events, similar to work by Nelson et al. (2006) and Yu and Riedl (2012). However, this representation significantly increases the size of the state space, and hence the quantity of training data that must be collected to adequately sample the space of narrative-centered tutorial planning policies. Consequently, this type of adjustment may be best studied under conditions where training data can be collected at significantly lower cost (e.g., web-based deployments not requiring active presence by researchers), or using simulated students where generating training data is not a bottleneck. Alternatively, automated feature selection techniques may provide insights into which subsets of plot points are most useful to model in narrative state representations.

We also did not find evidence of significant differences in student engagement between the experiment's two conditions. The most likely cause for this finding resembles the explanation for the content learning results. The Increase Urgency AES was one of the primary mechanisms included in CRYSTAL ISLAND to affect student engagement, in concert with several other AESs that were intended to cumulatively adjust the problem-solving scenario's overall difficulty. However, the Induced Planner rarely elected to increase the urgency in characters' dialogue, doing so during just 24% of opportunities. Moreover, the differences in difficulty between the two experimental conditions did not appear to be sufficient to affect students' self-reported presence or intrinsic motivation in CRYSTAL ISLAND.

## **8.2 Limitations**

To provide evidence in support of our thesis, this dissertation has presented empirical results that illustrate how the induced narrative-centered tutorial planner impacts students' problem-solving efficiency and problem-solving processes in narrative-centered learning environments. However, some of the key findings related to our research hypotheses were

supported by marginally significant evidence, or in some cases were not observed to be statistically significant (although they trended in the predicted direction). We attribute these results to insufficient statistical power for detecting the observed condition effects. Students demonstrated high levels of variance in problem-solving behaviors. Consequently, larger sample sizes may be necessary for future research examining students' problem-solving processes in CRYSTAL ISLAND.

In addition, there were some cases where our implemented narrative-centered tutorial planner did not take full advantage of all features in the concurrent Markov decision process framework for narrative-centered tutorial planning. Due to scheduling constraints related to working with public school systems, we were unable to implement tailored state representations and goal assignments for distinct AESs in the Induced Planner version of CRYSTAL ISLAND. Furthermore, we made limited use of arbitration procedures, implementing a domain-independent procedure for just one pair of AESs. These limitations do not diminish the findings we obtained on the effectiveness of the induced narrative-centered tutorial planner, which was designed and trained within the framework's requirements. In fact, we predict that tailored state representations, generated under the guidelines presented in Chapter 3, would yield even stronger results in the evaluation experiment with middle school students. Specifically, we expect that tailored state representations, obtained with automated feature selection, would yield narrative-centered tutorial planning policies that are better able to distinguish the states in which particular narrative adaptations enhance student outcomes. Future empirical work examining the impacts of tailored states, goal assignments, and domain-specific arbitration procedures will be valuable for further expanding on the findings reported in this dissertation.

In order to assess students' narrative-centered learning experiences, we employed several measures that have proven useful in prior research with CRYSTAL ISLAND. However, additional measures of student learning and engagement would be useful for providing a more comprehensive account of students' experiences in future experiments. In particular, this study did not include a measure of content knowledge transfer, which assesses students'

abilities to generalize content knowledge to new contexts. Additionally, we did not examine run-time measures of engagement, such as off-task behavior, to compliment presence and motivation questionnaire data. Prior collaborative research conducted by the author has demonstrated notable relationships between off-task behavior, learning and affective experiences (Rowe, McQuiggan, Robison, & Lester, 2009; Sabourin, Rowe, Mott, & Lester, in press). Empirical investigations of the induced narrative-centered tutorial planner's impact on students' off-task behaviors could help clarify how students experience disengagement and reengagement while solving the mystery.

This study also examined narrative-centered learning experiences that last for relatively short durations of time (approximately one hour). A key promise of narrative-centered learning environments is their capacity to enhance student motivation, which is likely to affect student learning by increasing time-on-task. Investigations of narrative-centered tutorial planners over extended periods of gameplay are an important direction for examining this work's capacity to scale to learning environments tailored for classroom use. Similarly, testing the framework on alternate narrative-centered learning environments, with distinct plotlines, would provide a valuable case study for demonstrating the framework's practical generalizability.

With regard to the limitations of the framework itself, one of the most significant restrictions is the assumption that AESs are fully independent of one another. In some ways, the independence assumption is a key contribution of the work; the concurrent Markov decision process formalization stems from recognition that a narrative-centered tutorial planning task can be factored in a similar way to other planning under uncertainty tasks, or joint probability distributions in general. The motivation for this assumption is analogous to the independence assumptions employed in Naïve Bayes classifiers, where strong independence assumptions are utilized to produce simple models that are practically useful, even in applications where the independence assumptions may not hold. However, we have not rigorously inspected under what conditions these independence assumptions actually hold in narrative-centered learning environments, nor have we examined how to recognize them

or to what extent they generalize to other interactive narratives. While the general premise of leveraging independence relationships within a narrative-centered tutorial planning task is promising, considerable work remains to examine how and when independence arises between AESs.

In this chapter, we have discussed results from an experiment evaluating the effectiveness of a narrative-centered tutorial planner implemented through our concurrent Markov decision process framework. A closer examination of the narrative-centered tutorial planner's constituent AESs has suggested that the Bryce's Revelation AES and Initial Test Count AES are the two most likely explanations for the observed differences in students' problem-solving efficiency and problem-solving processes between experimental conditions. Further, we speculate that the absence of observed condition effects on student learning are at least in part attributable to the particular narrative state representation we employed, which was driven by concerns about data sparsity. While limitations exist in both our experimental evaluation and planning framework, we have nevertheless provided a case for the effectiveness of our framework for inducing narrative-centered tutorial planning policies that tailor a range of narrative adaptations for enhancing student problem solving.

## CHAPTER 9

### Conclusion

The work described in this dissertation represents a notable step toward realizing the vision of personalized learning experiences that are simultaneously effective and engaging for students. We have presented a data-driven framework for dynamically tailoring story events in narrative-centered learning environments with concurrent Markov decision processes. Narrative-centered learning environments provide engaging learning experiences embedded within meaningful story contexts. They integrate salient features of narratives with digital game environments in order to increase student motivation, promote situated learning, and guide complex problem solving. Supporting students' learning processes in narrative-centered learning environments demands tailored assistance that is naturalistically, and discreetly, embedded within educational interactive narratives. Narrative-centered tutorial planners solve this problem by intelligently augmenting interactive story events in order to scaffold problem solving, learning, and engagement. Similar to effective one-on-one human tutoring, dynamically personalizing story events promotes tailored learning experiences that adapt pedagogical assistance according to the needs and preference of individual students.

Modular reinforcement learning provides a principled framework for devising effective narrative-centered tutorial planning policies that dynamically tailor story events. By modeling narrative-centered tutorial planning with concurrent Markov decision processes, we developed an effective narrative-centered tutorial planner for the CRYSTAL ISLAND learning environment. The implemented model performs a broad range of narrative adaptations dynamically to enhance students' problem-solving processes. After integrating the narrative-centered tutorial planner with CRYSTAL ISLAND, a controlled experiment with middle school students yielded empirical evidence in support of the following thesis: *modeling narrative-centered tutorial planning with concurrent Markov decision processes is an effective*

*approach for dynamically tailoring story events to enhance student problem solving in narrative-centered learning environments.*

## 9.1 Hypotheses Revisited

The research in this dissertation tested four hypotheses concerning the effectiveness of our framework for narrative-centered tutorial planning with concurrent Markov decision processes. The evaluation experiment, conducted with middle school students, produced the following results:

- **Hypothesis 1:** Students who use a narrative-centered learning environment with an Induced Planner *solve complex problem-solving tasks more quickly and more frequently* than students who use a baseline version of the environment.
  - Results showed that students in the Induced Planner condition solved CRYSTAL ISLAND's science problem-solving task in less time than students in the Baseline Planner condition. Furthermore, students in the Induced Planner condition trended toward solving CRYSTAL ISLAND's problem-solving task more frequently than students who interacted with the Baseline Planner version, or any comparable non-adaptive version of CRYSTAL ISLAND.
- **Hypothesis 2:** Students who use a narrative-centered learning environment with an Induced Planner *demonstrate more deliberate problem-solving practices* than students who use a baseline version of the environment.
  - Results from the evaluation experiment indicated that students in the Induced Planner group demonstrated greater efficiency during hypothesis testing. Specifically, students in the Induced Planner condition conducted significantly fewer tests after successfully identifying the transmission source in the laboratory than students in the Baseline Planner condition. Furthermore, students who used the Induced Planner version of CRYSTAL ISLAND demonstrated increased information-gathering behaviors prior to their hypothesis-testing phase of problem solving. Specifically, the Induced Planner

group demonstrated greater duration and frequency at conversing with virtual scientists about virus and bacteria concepts, as well as greater duration and frequency at viewing disease-focused posters about candidate diagnoses for the illness.

- **Hypothesis 3:** Students who use a narrative-centered learning environment with an Induced Planner *achieve greater content learning gains* than students who use a baseline version of the environment.
  - Students, regardless of the version of CRYSTAL ISLAND that they used, achieved significant gains in microbiology content knowledge. However, we did not find evidence that the Induced Planner produced greater content learning gains than the Baseline Planner. We anticipate that this negative result was due in part to the narrative state representation chosen for the Induced Planner. The state representation deviated from the Markov assumption; it emphasized plot states rather than plot trajectories. While the state representation was selected to avoid data sparsity issues during model training, it may have had limited ability to identify story trajectories that suggested a need for scaffolding content learning processes.
- **Hypothesis 4:** Students who use a narrative-centered learning environment with an Induced Planner *demonstrate greater engagement* than students who use a baseline version of the environment.
  - We did not find evidence that the Induced Planner produced greater student engagement than the Baseline Planner, as measured by presence and intrinsic motivation questionnaires. We infer that this negative result was a consequence of implementing narrative adaptations with too limited capacity to shape students' perceptions of presence and intrinsic motivation in CRYSTAL ISLAND. The immersive, aesthetic nature of CRYSTAL ISLAND's story world may have overpowered the Induced Planner's ability to affect student presence and intrinsic motivation. Similar to the learning results, we expect

that the chosen narrative state representation may have also impacted the framework's ability to induce policies that effectively enhance student engagement.

## 9.2 Summary

This dissertation has presented an empirical account of student learning, problem solving, and engagement in adaptive and non-adaptive versions of a narrative-centered learning environment CRYSTAL ISLAND. Analyses of middle school students' interactions with a non-adaptive version of CRYSTAL ISLAND found that the environment engenders significant learning gains, and that engagement in the narrative-centered learning environment is an important contributor to student learning and problem solving. Furthermore, an investigation of students' individual differences in CRYSTAL ISLAND indicated that high-achieving and low-achieving science students vary significantly in their gameplay behaviors, learning outcomes, and problem-solving outcomes. Specifically, high-achieving students utilize more traditional science resources within the narrative environment, such as virtual books and worksheets, as well as demonstrate higher presence and microbiology post-test scores. Low-achieving students gravitate toward novel gameplay features, such as virtual characters and lab equipment. These findings motivated the presented framework for devising narrative-centered tutorial planners that dynamically tailor story events in narrative-centered learning environments to scaffold students' learning, engagement, and problem-solving processes.

We have presented a novel framework for narrative-centered tutorial planning that involves decomposing an interactive narrative in terms of multiple independent sub-problems, formalized as adaptable event sequences. AESs provide a natural formalism for conceptualizing the sub-tasks of narrative-centered tutorial planning, and they readily describe several categories of narrative adaptations, including plot adaptations, discourse adaptations, and user tailoring. After decomposing the interactive narrative, AESs are modeled as concurrent Markov decision processes, with rewards based on students' experiential outcomes such as learning gains and problem-solving results. Policies for

solving the MDPs are obtained through reinforcement learning techniques, such as certainty equivalent learning. Training data for learning the narrative adaptation policies is collected during human subject studies with an exploratory version of the narrative-centered tutorial planner. Arbitration techniques are employed to resolve conflicts between multiple competing AESs. Within the outlines of this framework, we have demonstrated that tailoring the state representations used to encode MDPs in a narrative-centered tutorial planner dramatically reduces training data requirements compared to an analogous centralized planner. Reductions in training data requirements are especially valuable when learning models directly from student data, which is resource-intensive to obtain.

The framework's effectiveness was empirically demonstrated in an evaluation experiment involving middle school students using an implemented narrative-centered tutorial planner for the CRYSTAL ISLAND learning environment. This is the first deployed narrative-centered tutorial planner to have been directly induced from student interaction data; the model's narrative adaptation behaviors did not require hand authoring or machine learning from human demonstrations. In order to induce this model, a sizable training corpus was generated through a series of studies involving middle school students interacting with a version of CRYSTAL ISLAND imbued with an exploratory narrative-centered tutorial planner. Using this training corpus, we successfully induced a composite narrative-centered tutorial planner comprised of 12 narrative adaptation policies, one for each AES in the adaptive version of CRYSTAL ISLAND. The resulting planner satisfied the four operational requirements of narrative-centered learning environments: 1) performing instructional interventions that are naturally embedded within an interactive narrative, 2) accommodating a diverse range of potential story manipulations, 3) explicitly addressing the instructional goals of narrative-centered learning, and 4) satisfying the real-time performance requirements of interactive narratives.

The evaluation experiment with middle school students revealed that the Induced Planner significantly improved students' efficiency and deliberateness in problem solving compared to a baseline system. Specifically, students solved the mystery in less time, and

more often, when interacting with the Induced Planner, and they demonstrated improved information gathering and hypothesis testing behaviors. In addition to demonstrating the effectiveness of our modular reinforcement learning framework for narrative-centered tutorial planning, we presented an empirical account of students' gameplay behaviors—such as conversing with virtual characters, reading in-game books, viewing in-game posters, running tests in a virtual laboratory, and completing a diagnosis worksheet—in the adaptive version of CRYSTAL ISLAND. The results provide substantial evidence in support of our thesis that modeling narrative-centered tutorial planning with concurrent Markov decision processes is an effective approach for dynamically tailoring story events to enhance student problem solving in narrative-centered learning environments.

### **9.3 Future Work**

While this dissertation has presented considerable evidence to support its thesis, several directions remain for future work. First, additional empirical investigations of the narrative-centered tutorial planning framework's salient features—such as tailoring state representations to individual AESs, mapping AESs to multiple goals, implementing domain-dependent arbitration procedures, and pursuing alternate reinforcement learning techniques—will provide further insights into the framework's strengths and weaknesses. In the future, we plan to begin this work by using existing data, examining how alternate state configurations impact narrative adaptation policies induced through our framework. Furthermore, we expect to directly examine the impacts of data sparsity in training corpora by investigating the stability of narrative adaptation policies induced with varying quantities of training data. We are also poised to investigate on-line reinforcement learning techniques for inducing narrative adaptation policies. To do this, we will use the existing training corpus to model simulated students for dynamically generating training episodes. Now that we have demonstrated the practical effectiveness of narrative-centered tutorial planning policies induced by off-line reinforcement learning techniques, we can begin to compare policies induced by on-line and off-line methods.

Second, we intend to further examine the framework's theoretical properties, extending our analyses beyond training data requirements and time complexity as presented in this dissertation. One of the motivations for decomposing narrative-centered tutorial planning tasks into multiple independent AESs is the belief that decomposed planners will perform comparably to centralized planners while requiring less training data. A formal evaluation of this postulate would compliment our existing results, clarify the extent to which the two forms of models are similar, and reveal the tradeoffs between the two approaches. A related component of our framework is the assumption that distinct AESs are independent. This simplifying assumption also merits further examination to identify the conditions under which it holds and determine how to recognize independence relationships within narrative-centered tutorial planning tasks.

Looking beyond the computational properties of our framework, future work will include additional studies involving human participants interacting with the implemented narrative-centered tutorial planner. In order to obtain more comprehensive accounts of student learning, assessments of knowledge transfer will prove essential. Instilling transferrable knowledge is one of the key goals of education, and it provides the basis for most educational assessments. Supplemental analyses of students' problem-solving processes using educational data mining techniques, such as differential sequence mining, will also likely yield insights on the implemented narrative-centered tutorial planner's impacts (Kinnebrew and Biswas, 2012). The statistical analyses in this dissertation have primarily considered aggregate measures of students' problem-solving behaviors. Educational data mining techniques provide alternate methods for analyzing students' problem-solving behavior that preserve the sequential orders of students' actions. In order to further examine student engagement, analyses of students' off-task behaviors are likely to prove informative. Students' off-task behaviors reveal patterns of disengagement and reengagement, providing objective measures of student engagement that compliment the questionnaires that we have already described.

Finally, it will be critical to test the generalizability of our framework by implementing narrative-centered tutorial planners in alternate narrative-centered learning environments. We have sought to maintain the framework's generalizability through the use of Markov decision processes—which provide a general mathematical framework for sequential decision-making—but a demonstrated implementation of a narrative-centered tutorial planner for a distinct learning environment would substantially strengthen these arguments. Fortunately, CRYSTAL ISLAND is part of a family of narrative-centered learning environments for K-12 science subjects developed in our laboratory, each with distinct narratives, curricula, and target audiences. Implementing narrative-centered tutorial planners for these alternate environments, using the procedures specified by our framework, is feasible, and it will provide a practical test of the approach's applicability beyond a single context. In particular, implementing narrative-centered tutorial planners for environments intended to last for multiple class periods will be essential for demonstrating the framework's applicability to classroom settings, which often demand educational resources that can support routine use and long-term educational impact.

#### **9.4 Concluding Remarks**

Over the past few decades, several indicators have signaled that the United States has developed an emerging gap in STEM interest and achievement that causes many bright young men and women to fail to reach their potential (President's Council 2010; National Science Board 2010). While the nature of these problems is complex, computational advances in learning technologies can serve a critical role in the solutions. This dissertation was inspired by a desire to facilitate efforts to integrate student engagement, as fostered by modern commercial games, with deep learning experiences, as promoted by effective one-on-one human tutoring. Narrative-centered learning environments create opportunities for situated learning experiences that would have been unimaginable for prior generations of students. By leveraging features of intelligent tutoring systems, narrative-centered tutorial planners support the promise of personalized instruction for every student. Given the gravity

of today's educational challenges, it is with a sense of urgency that the research community should pursue research in game-based learning, intelligent tutoring systems, and narrative-centered tutorial planning.

## REFERENCES

- Abbeel, P., Coates, A., Quigley, M., & Ng, A. Y. (2007). An application of reinforcement learning to aerobatic helicopter flight. *Proceedings of the 21st Annual Conference on Neural Information Processing Systems*, Vancouver, Canada. 1–8.
- Amato, C., & Shani, G. (2010). High-level reinforcement learning in strategy games. *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, Toronto, Canada. 75–82.
- Anderson, J. R., Boyle, C. F., Corbett, Albert T., & Lewis, M. W. (1990). Cognitive Modeling and Intelligent Tutoring. *Artificial intelligence and learning environments*, 42(1), 7–49.
- Aylett, R., & Louchart, S. (2008). If I were you - Double appraisal in affective agents. *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal. 1233–1236.
- Aylett, R., Louchart, S., Dias, J., Paiva, A., & Vala, M. (2005). FearNot! An experiment in emergent narrative. *Proceedings of the 5th International Conference on Intelligent Virtual Agents*, Kos, Greece. 305–316.
- Barab, S. A., Gresalfi, M., & Ingram-Goble, A. A. (2010). Transformational play: Using games to position person, content, and context. *Educational Researcher*, 39(7), 525–536.
- Barber, H., & Kudenko, D. (2007). Dynamic generation of dilemma-based interactive narratives. *Proceedings of the 3rd Artificial Intelligence and Interactive Digital Entertainment Conference*. Palo Alto, CA. 2–7.

- Barnes, T., & Stamper, J. (2008). Toward automatic hint generation for logic proof tutoring using historical student data. *Proceedings of the 9th International Conference on Intelligent Tutoring Systems*, Montreal, Canada. 373–382.
- Barto, A.G., & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4), 341–379.
- Beck, J. E., Woolf, B. P., & Beal, C. R. (2000). ADVISOR : A machine learning architecture for intelligent tutor construction. *Proceedings of the 17th National Conference on Artificial Intelligence*, Austin, Texas. 552–557.
- Bhat, S., Isbell, C. L., & Mateas, M. (2006). On the difficulty of modular reinforcement learning for real-world partial programming. *Proceedings of the 21st National Conference on Artificial Intelligence*. Boston, Massachusetts. 318–323.
- Boyer, K. E., Phillips, R., Ingram, A., Ha, E. Y., Wallis, M., Vouk, M., et al. (2010). Characterizing the effectiveness of tutorial dialogue with hidden Markov models. *Proceedings of the 10th International Conference on Intelligent Tutoring Systems*, Pittsburgh, PA. 55–64.
- Britner, S. L., & Pajares, F. (2006). Sources of science self-efficacy beliefs of middle school students. *Journal of Research in Science Teaching*, 43(5), 485–499.
- Britt, M. A., Wiemer-Hastings, P., Larson, A. A., & Perfetti, C. A. (2004). Using intelligent feedback to improve sourcing and integration in students' essays. *International Journal of Artificial Intelligence in Education*, 14(3), 359-374.
- Bruner, J. (1991). The narrative construction of reality. *Critical Inquiry*, 18(1), 1-21.
- Cavazza, M., Charles, F., & Mead, S. J. (2002). Character-based interactive storytelling. *IEEE Intelligent Systems*, 17(4), 17–24.

- Chi, M., Jordan, P., VanLehn, K., & Hall, M. (2008). Reinforcement learning-based feature selection for developing pedagogically effective tutorial dialogue tactics. *Proceedings of the 1st International Conference on Educational Data Mining*, Montreal, Canada. 258–265.
- Chi, M., VanLehn, K., & Litman, D. (2010). Do micro-level tutorial decisions matter: Applying reinforcement learning to induce pedagogical tutorial tactics. *Proceedings of the 10th International Conference on Intelligent Tutoring Systems*, Pittsburgh, PA. 224–234.
- Corbett, A T, & Anderson, J R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4), 253-278.
- Doya, K. (2000). Reinforcement learning in continuous time and space. *Neural Computation*, 12(1), 219-45.
- Elliot, A., & McGregor, H. (2001). A 2x2 achievement goal framework. *Journal of Personality and Social Psychology*, 80(3), 501–519.
- Fairclough, C. R. (2004). Story games and the OPIATE system. Ph. D. Thesis, University of Dublin – Trinity College Department of Computer Science.
- Federation of American Scientists. (2006). Harnessing the power of video games for learning. Summit on Educational Games. Retrieved from <http://fas.org/gamesummit/>
- Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4), 189–208.
- Garnefski, N., & Kraaij, V. (2006). Cognitive emotion regulation questionnaire: Development of a short 18-item version. *Personality and Individual Differences*, 41(6), 1045–1053.

- Gee, J. P. (2007). *What video games have to teach us about learning and literacy* (2<sup>nd</sup> ed.). New York: Palgrave Macmillan.
- Graesser, A. C., Lu, S., Jackson, G. T., Mitchell, H. H., Ventura, M., Olney, A., et al. (2004). AutoTutor: A tutor with dialogue in natural language. *Behavior Research Methods, Instruments, & Computers*, 36(2), 180-192.
- Graesser, A. C., & Ottati, V. (1996). Why stories? Some evidence, questions, and challenges. In R. S. Wyer (Ed.), *Knowledge and memory: The real story* (pp. 121-132). Hillsdale, NJ: Erlbaum.
- Hickey, D. T., Ingram-Goble, A. A., & Jameson, E. M. (2009). Designing assessments and assessing designs in virtual educational environments. *Journal of Science Education and Technology*, 18(2), 187-208.
- Hmelo-Silver, C. E., Duncan, R. G., & Chinn, C. A. (2007). Scaffolding and achievement in problem-based and inquiry learning: A response to Kirschner, Sweller, and Clark (2006). *Educational Psychologist*, 42(2), 99-107.
- Honey, M. A., & Hilton, M. (2011). Learning science through computer games and simulations. Retrieved from [http://www.nap.edu/catalog.php?record\\_id=13078](http://www.nap.edu/catalog.php?record_id=13078)
- Jaakkola, T., Singh, S. P., & Jordan, M. I. (1995). Reinforcement learning algorithm for partially observable Markov decision problems. *Proceedings of the 10th Annual Conference on Advances in Neural Information Processing Systems*, Denver, Colorado. 345-352.
- Jeong, H., Gupta, A., Roscoe, R., Wagster, J., & Biswas, G. (2008). Using hidden Markov models to characterize student behaviors in learning-by-teaching environments. *Proceedings of the 9th International Conference on Intelligent Tutoring Systems*, Montreal, Canada. 614-625.

- John, O. P., Naumann, L. P., & Soto, C. J. (2008). Paradigm shift to the integrative Big Five trait taxonomy: History, measurement, and conceptual issues. In O. P. John, R. W. Robins, & L. A. Pervin (Eds.), *Handbook of personality: Theory and research* (pp. 114–158). New York, NY: Guilford Press.
- Johnson, W. L. (2010). Serious use of a serious game for language learning. *International Journal of Artificial Intelligence in Education*, 20(2), 175–195.
- Jonassen, D. H., & Hernandez-Serrano, J. (2002). Case-based reasoning and instructional design: Using stories to support problem solving. *Educational Technology Research and Development*, 50(2), 65–77.
- Jonassen, D. H. (2000). Toward a design theory of problem solving. *Educational Technology Research and Development*, 48(4), 63–85.
- Kaelbling, L. P., Littman, Michael L, & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
- Karlsson, J. (1997). Learning to solve multiple goals. Ph. D. Thesis, University of Rochester Department of Computer Science.
- Ketelhut, D. J., Dede, C., Clarke, J., & Nelson, B. C. (2010). A multi-user virtual environment for building higher order inquiry skills in science. *British Journal of Educational Technology*, 41(1), 56–68.
- Kim, J., Hill, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., et al. (2009). BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3), 289–308.
- Kinnebrew, J. S., & Biswas, G. (2012). Identifying learning behaviors by contextualizing differential sequence mining with action features and performance evolution.

*Proceedings of the Fifth International Conference on Educational Data Mining*, Chania, Greece. 57–64.

- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist*, 41(2), 75–86.
- Koedinger, K. R., Anderson, J.R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8(1), 30–43.
- Lee, S., Mott, B. W., & Lester, J. C. (2010). Optimizing story-based learning: An investigation of student narrative profiles. *Proceedings of the 10th International Conference on Intelligent Tutoring Systems*, Pittsburgh, PA. 155–165.
- Lee, S. Y., Mott, B. W., & Lester, J. C. (2012). Real-time narrative-centered tutorial planning for story-based learning. *Proceedings of the 11th International Conference on Intelligent Tutoring Systems*, Chania, Greece. 476–481.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. *Proceedings of the 11th International Conference on Machine Learning*, New Brunswick, NJ. 157-163.
- Littman, M. L., Dean, T. L., & Kaelbling, L. P. (1995). On the complexity of solving Markov decision problems. *Proceedings of the 11<sup>th</sup> Conference on Uncertainty in Artificial Intelligence*, Montreal, Canada. 394–402.
- Marsella, S., Johnson, W. L., & LaBore, C. (2000). Interactive pedagogical drama. *Proceedings of the Fourth International Conference on Autonomous Agents*, Barcelona, Spain. 301–308.

- Marsella, S., Johnson, W. L., & Labore, C. (2003). Interactive pedagogical drama for health interventions. *Proceedings of the 11th International Conference on Artificial Intelligence in Education*. Sydney, Australia. 341–348.
- Mateas, M., & Stern, A. (2005). Structuring content in the Façade interactive drama architecture. *Proceedings of 1st International Conference on Artificial Intelligence and Interactive Digital Entertainment*, Marina del Rey, CA. 93–98.
- McAuley, E., Duncan, T., & Tammen, V. E. (1989). Psychometric properties of the Intrinsic Motivation Inventory in a competitive sport setting: A confirmatory factor analysis. *Research Quarterly for Exercise and Sport*, 60(1), 48–58.
- McQuiggan, S. W., Mott, B. W., & Lester, J. C. (2008). Modeling self-efficacy in intelligent tutoring systems: An inductive approach. *User Modeling and User-Adapted Interaction*, 18(1-2), 81–123.
- Mislevy, R. J., & Gitomer, D. H. (1995). The role of probability-based inference in an intelligent tutoring system. *User Modeling and User-Adapted Interaction*, 5(3), 253–282.
- Mitrovic, A., & Ohlsson, S. (1999). Evaluation of a constraint-based tutor for a database language. *International Journal of Artificial Intelligence in Education*, 10(3-4), 238–256.
- Mott, B. W., Lee, S., & Lester, J. C. (2006). Probabilistic goal recognition in interactive narrative environments. *Proceedings of the 21st National Conference on Artificial Intelligence*, Boston, MA. 187–192.
- Mott, B. W., & Lester, J. C. (2006a). U-Director : A decision-theoretic narrative planning architecture for storytelling environments. *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*, Hakodate, Japan. 977–984.

- Mott, B. W., & Lester, J. C. (2006b). Narrative-centered tutorial planning for inquiry-based learning environments. *Proceedings of the Eighth International Conference on Intelligent Tutoring Systems*, Jhongli, Taiwan. 675–684.
- Murray, J. H. (1998). *Hamlet on the holodeck: The future of narrative in cyberspace*. Cambridge, MA: MIT Press.
- Murray, R. C., Vanlehn, K., & Mostow, J. (2004). Looking ahead to select tutorial actions: A decision-theoretic approach. *International Journal of Artificial Intelligence in Education*, 14(3-4), 235–278.
- National Science Board (2010). *Preparing the Next Generation of STEM Innovators: Identifying and Developing our Nation's Human Capital* (NSF Publication No. NSB-10-33). Arlington, VA: National Science Board. Retrieved from <http://www.nsf.gov/nsb/publications/2010/nsb1033.pdf>
- Nelson, B. C. (2007). Exploring the use of individualized, reflective guidance in an educational multi-user virtual environment. *Journal of Science Education and Technology*, 16(1), 83–97.
- Nelson, B. C., & Ketelhut, D. J. (2008). Exploring embedded guidance and self-efficacy in educational multi-user virtual environments. *International Journal of Computer-Supported Collaborative Learning*, 3(4), 413–427.
- Nelson, M. J., & Mateas, M. (2005). Search-based drama management in the interactive fiction Anchorhead. *Proceedings of the First Annual Conference on Artificial Intelligence and Interactive Digital Entertainment*, Marina del Rey, CA. 99–104.
- Nelson, M. J., Roberts, D. L., Isbell, C. L., & Mateas, M. (2006). Reinforcement learning for declarative optimization-based drama management. *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*, Hakodate, Japan. 775–782.

- Nietfeld, J. L., Cao, L., & Osborne, J. W. (2006). The effect of distributed monitoring exercises and feedback on performance, monitoring accuracy, and self-efficacy. *Metacognition and Learning, 1*(2), 159–179.
- Polkinghorne, D. (1988). *Narrative knowing and the human sciences*. Albany, NY: State University of New York Press.
- Porteous, J., Cavazza, M., & Charles, F. (2010). Narrative generation through characters' point of view. *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, Toronto, Canada. 1297–1304.
- Prensky, M. (2001). *Digital game-based learning*. New York: McGraw-Hill.
- President's Council of Advisors on Science and Technology. (2010). *Prepare and Inspire: K-12 Education in Science, Technology, Engineering, and Math (STEM) for America's Future*. Washington, D.C.: Executive Office of the President. Retrieved from <http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-stem-ed-final.pdf>
- Pynadath, D., & Marsella, S. (2005). Psychsim: Modeling theory of mind with decision-theoretic agents. *Proceedings of the 19th International Joint Conference on Artificial Intelligence*. Edinburgh, Scotland, UK. 1181–1186.
- Rickel, J., & Johnson, W. L. (1999). Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence, 13*(4), 343–382.
- Riedl, M., Saretto, C. J., & Young, R. M. (2003). Managing interaction between users and agents in a multi-agent storytelling environment. *Proceedings of the 2nd International Conference on Autonomous Agents and Multiagent Systems*, Melbourne, Australia. 741–748.

- Riedl, M., Stern, A., Dini, D. M., & Alderman, J. M. (2008). Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning*, 4(2), 23–42.
- Riedmiller, M., Gabel, T., Hafner, R., & Lange, S. (2009). Reinforcement learning for robot soccer. *Autonomous Robots*, 27(1), 55–73.
- Roberts, D. L., & Isbell, C. L. (2008). A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning*, 4(2), 61–75.
- Roberts, D. L., Nelson, M. J., Isbell, C. L., Mateas, M., & Littman, M. L. (2006). Targeting specific distributions of trajectories in MDPs. *Proceedings of the 21st National Conference on Artificial Intelligence*, Boston, MA. 1213–1218.
- Robison, J., McQuiggan, S., & Lester, J. (2009). Modeling task-based vs. affect-based feedback behavior in pedagogical agents: An inductive approach. *Proceedings of the 14th International Conference on Artificial Intelligence in Education*. Brighton, UK. 25–32.
- Rowe, J., Ha, E. Y., & Lester, J. (2008). Archetype-driven character dialogue generation for interactive narrative. *Proceedings of the 8th International Conference on Intelligent Virtual Agents*, Tokyo, Japan. 45–58.
- Rowe, J., & Lester, J. (2010). Modeling user knowledge with dynamic Bayesian networks in interactive narrative environments. *Proceedings of the 6th Annual Artificial Intelligence and Interactive Digital Entertainment Conference*. Palo Alto, CA. 57–62.

- Rowe, J., McQuiggan, S., Robison, J., & Lester, J. (2009). Off-task behavior in narrative-centered learning environments. *Proceedings of the 14th International Conference on Artificial Intelligence in Education*, Brighton, UK. 99–106.
- Rowe, J., Shores, L., Mott, B., & Lester, J. (2010a). Individual differences in gameplay and learning: A narrative-centered learning perspective. *Proceedings of the 5th International Conference on Foundations of Digital Games*, Monterey, CA. 171–178.
- Rowe, J., Shores, L., Mott, B., & Lester, J. (2010b). A framework for narrative adaptation in interactive story-based learning environments. *Proceedings of the FDG '10 Workshop on Intelligent Narrative Technologies III*. Monterey, CA.
- Rowe, J., Shores, L., Mott, B., & Lester, J. (2011). Integrating learning, problem solving, and engagement in narrative-centered learning environments. *International Journal of Artificial Intelligence in Education*, 21(1), 1–19.
- Ryan, R. M. (1982). Control and information in the intrapersonal sphere: An extension of Cognitive Evaluation Theory. *Journal of Personality and Social Psychology*, 43(3), 450–461.
- Sabourin, J., Mott, B., & Lester, J. (2011). Modeling learner affect with theoretically grounded dynamic Bayesian networks. *Proceedings of the 4<sup>th</sup> International Conference on Affective Computing and Intelligent Interaction*, Memphis, TN. 286–295.
- Sabourin, J., Rowe, J., Mott, B., & Lester, J. (2012). Exploring inquiry-based problem-solving strategies in game-based learning environments. *Proceedings of the 11th International Conference on Intelligent Tutoring Systems*, Chania, Greece. 470–475.
- Sabourin, J., Rowe, J. P., Mott, B. W., & Lester, J. C. (in press). Considering alternate features to classify off-task behavior as emotion self-regulation: A supervised learning

approach. To appear in *Journal of Educational Data Mining*.

Sabourin, J., Rowe, J., Mott, B., & Lester, J. (2011). When off-task is on-task: The affective role of off-task behavior in narrative-centered learning environments. *Proceedings of the Fifteenth International Conference on Artificial Intelligence in Education*, Auckland, New Zealand. 534–536.

Schraw, G. (1997). Situational interest in literary text. *Contemporary Educational Psychology*, 22(4), 436–456.

Schubert, T., Friedmann, F., & Regenbrecht, H. (1999). Embodied presence in virtual environments. In R. Paton (Ed.), *Visual representations and interpretations* (pp. 269–278). London: Springer-Verlag.

Schwartz, D., & Bransford, J. (1998). A time for telling. *Cognition and Instruction*, 16(4), 475–522.

Shaffer, D. W. (2006). *How computer games help children learn*. New York: Palgrave Macmillan.

Sharma, M., Holmes, M., Santamaria, J., Irani, A., Isbell, C. L., & Ram, A. (2007). Transfer learning in real-time strategy games using hybrid CBR/RL, *Proceedings of the 20<sup>th</sup> International Joint Conference on Artificial Intelligence*, Hyderabad, India. 1041–1046.

Sharma, M., Ontanón, S., Strong, C., Mehta, M., & Ram, A. (2007). Towards player preference modeling for drama management in interactive stories. *Proceedings of the 20<sup>th</sup> International Florida Artificial Intelligence Research Society Conference*, Key West, Florida. 571–576.

- Shores, L. R., Rowe, J. P., & Lester, J. C. (2011). Early prediction of cognitive tool use in narrative-centered learning environments. *Proceedings of the 15<sup>th</sup> International Conference on Artificial Intelligence in Education*, Auckland, New Zealand. 320–327.
- Shute, V. J. (1995). SMART: Student modeling approach for responsive tutoring. *User Modeling and User-Adapted Interaction*, 5(1), 1–44.
- Shute, V. J., Masduki, I., Donmez, O., Dennen, V. P., Kim, Y.-jeon, Jeong, A. C., et al. (2010). Modeling, assessing, and supporting key competencies within game environments. In D. Ifenthaler, P. Pirnay-Dummer, & N. M. Seel (Eds.), *Computer-Based Diagnostics and Systematic Analysis of Knowledge* (pp. 281–309). New York: Springer-Verlag.
- Shute, V. J. & Psotka, J. (1996). Intelligent tutoring systems: Past, present and future. In D. Jonassen (Ed.), *Handbook of Research on Educational Communications and Technology*. Scholastic.
- Si, M., Marsella, S., & Pynadath, D. (2005). Thespian: An architecture for interactive pedagogical drama. *Proceedings of the 12th International Conference on Artificial Intelligence in Education*. Amsterdam, the Netherlands. 595–602.
- Si, M., Marsella, S., & Pynadath, D. (2009). Directorial control in a decision-theoretic framework for interactive narrative. *Proceedings of the 2nd International Conference on Interactive Digital Storytelling*, Guimaraes, Portugal. 221–233.
- Si, M., Marsella, S., & Riedl, M. (2008). Integrating story-centric and character-centric processes for authoring interactive drama. *Proceedings of the 4th Artificial Intelligence and Interactive Digital Entertainment Conference*, Palo Alto, CA. 203–208.

- Singh, S., Kearns, M., Litman, D., & Walker, M. (1999). Reinforcement learning for spoken dialogue systems. *Proceedings of the 13th Annual Conference on Neural Information Processing Systems*, Denver, CO. 956–962.
- Singh, S., Litman, D., Kearns, M., & Walker, M. (2002). Optimizing dialogue management with reinforcement learning: Experiments with the NJFun system. *Journal of Artificial Intelligence Research*, 16(1), 105–133.
- Spires, H., Rowe, J., Mott, B., & Lester, J. (2011). Problem solving and game-based learning: Effects of middle grade students' hypothesis testing strategies on science learning outcomes. *Journal of Educational Computing Research*, 44(4), 453–472.
- Sprague, N., & Ballard, D. (2003). Multiple-goal reinforcement learning with modular sarsa (0). *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, Acapulco, Mexico. 1445–1447.
- Stone, P., Sutton, R. S., & Kuhlmann, G. (2005). Reinforcement learning for RoboCup-Soccer keepaway. *Adaptive Behavior*, 13(3), 165-188.
- Sutton, R. S., & Barto, Andrew G. (1998). *Reinforcement learning: An Introduction*. Cambridge, MA: MIT Press.
- Swanson, R., & Gordon, A. (2008). Say Anything: A massively collaborative open domain story writing companion. *Proceedings of the 1st Joint International Conference on Interactive Digital Storytelling*, Erfurt, Germany. 32–40.
- Tetreault, J. R., Bohus, D., & Litman, D. J. (2007). Estimating the reliability of MDP policies: A confidence interval approach. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies*, Rochester, NY. 276–283.

- Tetreault, J. R., & Litman, D. J. (2006). Comparing the utility of state features in spoken dialogue using reinforcement learning. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies*, New York, NY. 272–279.
- Thomas, J. (2011). Automated scaffolding of task-based learning in non-linear game environments. Ph.D. Thesis, North Carolina State University Department of Computer Science.
- Thomas, J. M., & Young, R.M. (2010). Annie: Automated generation of adaptive learner guidance for fun serious games. *IEEE Transactions on Learning Technologies*, 3(4), 329–343.
- Thue, D., Bulitko, V., Spetch, M., & Wasylshen, E. (2007). Interactive storytelling: A player modelling approach. *Proceedings of the 3rd Artificial Intelligence and Interactive Digital Entertainment Conference*, Palo Alto, CA. 43–48.
- Traum, D., Marsella, S., Gratch, J., Lee, J., & Hartholt, A. (2008). Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. *Proceedings of the 8th International Conference on Intelligent Virtual Agents*, Tokyo, Japan. 117–130.
- VanLehn, K. (2006). The Behavior of Tutoring Systems. *International Journal of Artificial Intelligence in Education*, 16(3), 227–265.
- VanLehn, K., Lynch, C., Schulze, K., Shapiro, J. A., Shelby, R., Taylor, L., et al. (2005). The Andes physics tutoring system: Lessons learned. *International Journal of Artificial Intelligence in Education*, 15(3), 147–204.
- Vannini, N., Enz, S., Sapouna, M., Wolke, D., Watson, S., Woods, S., Dautenhahn, K., et al. (2011). “FearNot!”: A computer-based anti-bullying-programme designed to foster peer intervention. *European Journal of Psychology of Education*, 26(1), 21–44.

- Warren, S. J., Dondlinger, M. J., & Barab, S. A. (2008). A MUVE towards PBL writing: Effects of a digital learning environment designed to improve elementary student writing. *Journal of Research on Technology in Education*, 41(1), 113–140.
- Watson, S., Vannini, N., Woods, S., Dautenhahn, K., Sapouna, M., Enz, S., et al. (2010). Inter-cultural differences in response to a computer-based anti-bullying intervention. *Educational Research*, 52(1), 61–80.
- Weyhrauch, P. (1997). Guiding Interactive Drama. Ph. D. Thesis, Carnegie Mellon University Department of Computer Science.
- Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3), 225–240.
- Woolf, B. P. (2008). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Burlington, MA: Morgan Kaufmann.
- Young, R. Michael. (1999). Notes on the use of plan structures in the creation of interactive plot. *Working Notes of the 1999 AAAI Fall Symposium on Narrative Intelligence*, North Falmouth, MA. 164–167.
- Yu, H., & Riedl, M. O. (2012). A sequential recommendation approach for interactive personalized story generation. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, Valencia, Spain. 71–78

## **APPENDICES**

## **Appendix A: Demographics and Game-Playing Experience Survey**

In the near future you will be playing an educational game called CRYSTAL ISLAND, which is being developed at North Carolina State University. The results from this survey will help us to improve the CRYSTAL ISLAND game.

The survey will take approximately thirty minutes to complete. Please answer the questions as honestly and thoroughly as you can. Your answers will be kept confidential. We greatly appreciate your time and input.

1. What is your full name?
2. What is your gender?
  - a.) Male
  - b.) Female
3. What is your age?
4. What is your race?
  - a.) American Indian or Alaska Native
  - b.) Asian
  - c.) Black or African American
  - d.) Hispanic or Latino
  - e.) Native Hawaiian or Other Pacific Islander
  - f.) Other (please specify)
5. Do you play video games?
  - a.) Not at all
  - b.) Rarely
  - c.) Occasionally
  - d.) Frequently
  - e.) Very frequently
6. How skilled are you when playing video games?

- a.) Not at all
  - b.) Limited skills
  - c.) Average
  - d.) Skilled
  - e.) Very skilled
7. How many hours per week do you typically play video games?

### **Appendix B: Self-Efficacy for Science Questionnaire**

This questionnaire is designed to help us gain a better understanding of the kinds of things that create difficulties for students in science classes. Please indicate your opinion about each of the following statements.

The following statement is \_\_\_\_\_ like me.

1 = Nothing, 2 = Very Little, 3 = Somewhat, 4 = Quite a Bit, 5 = A Great Deal

1. I am sure that I can learn science.
2. I can get a good grade in science.
3. I am sure I could do high school level work in science.
4. I have a lot of self-confidence when it comes to science.
5. I am not the type to do well in science.
6. It takes me a long time to catch on to new topics in science.
7. Even before I begin a new topic in science, I feel confident I'll be able to understand it.
8. I think I have good skills and strategies to learn science.

**Appendix C: CRYSTAL ISLAND Curriculum Test v2.0**

Mark the correct answer to each question.

1. Which of the following is true:
  - a.) Bacteria are not alive.
  - b.) Viruses are not alive.
  - c.) Viruses are alive.
  - d.) Both bacteria and viruses are alive.
2. Rank the order of bacteria, fungi, and viruses from the smallest to the largest.
  - a.) Fungi, Bacteria, Viruses
  - b.) Viruses, Bacteria, Fungi
  - c.) Bacteria, Viruses, Fungi
  - d.) Bacteria, Fungi, Viruses
3. You notice that the organism you are looking at through the microscope does not have a nucleus. What type of organism might you be looking at?
  - a.) A parasite or a carcinogen
  - b.) A fungus or a bacterium
  - c.) A bacterium or a virus
  - d.) A fungus or a parasite
4. Your friend is feeling ill and goes to the doctor. The doctor gives your friend an antibiotic and as a result your friend begins to feel better soon afterwards. What microorganism likely made your friend sick?
  - a.) A bacterium
  - b.) A virus
  - c.) Cancer
  - d.) A parasite

5. Your lab partners are examining a pathogen through a microscope and need your help identifying what type of pathogen they are looking at. They are unable to tell the size of the organism, but they can see that it is smooth and round in shape. What pathogen are your lab partners looking at?
  - a.) A virus
  - b.) A bacterium
  - c.) A fungus
  - d.) Both a virus and a bacterium
6. Which of the following statements about viruses is true?
  - a.) Viruses reproduce through binary fission.
  - b.) Viruses can be viewed through a light microscope.
  - c.) Viruses consist of genetic material within a capsid.
  - d.) Viruses are considered the smallest living cells.
7. A disease that is caused by bacteria is
  - a.) E-coli.
  - b.) Influenza.
  - c.) Smallpox.
  - d.) Ebola.
8. All of the following can be considered pathogens, EXCEPT
  - a.) Bacteria.
  - b.) Fungi.
  - c.) Viruses.
  - d.) Carcinogens.
9. You are a scientist trying to determine what type of microscope you should use to view a certain pathogen. To make your decision, you should consider all of the following EXCEPT:
  - a.) Type of pathogen.
  - b.) The size of the pathogen.

- c.) Whether the pathogen is living or dead.
  - d.) The amount of light in the room.
10. You are part of a team of scientists investigating an illness that may have been caused by contaminated food. Your team has asked the question, “What food source could be causing the illness?” What step should you take next?
- a.) You must perform an experiment.
  - b.) You must generate a hypothesis.
  - c.) You must gather information.
  - d.) You must make observations.
11. What pathogen is considered the smallest LIVING microorganism?
- a.) Bacteria
  - b.) Viruses
  - c.) Fungi
  - d.) Parasites
12. An illness has been spreading through a town, and doctors have told sick people to remain at home until they are well. The spread of the illness appears to have been stopped due to the doctors’ instructions. What does this tell you must be true about the illness?
- a.) It was caused by a mutagen because it spread from one organism to another.
  - b.) It was caused by a pathogen because it spread from one organism to another.
  - c.) It was caused by a carcinogen because it spread from one organism to another.
  - d.) It was not caused by a mutagen, carcinogen, or pathogen.
13. All of the following are characteristics of a pathogen EXCEPT:
- a.) Pathogens can spread.
  - b.) Pathogens only cause disease in living creatures.
  - c.) Pathogens can have more than one route of transmission.
  - d.) Pathogens are all alive.
14. You have determined that your patient’s disease has been caused by a genetic mutation. Knowing this you can determine that the disease was caused by:

- a.) A bacteria.
  - b.) A mutagen.
  - c.) A carcinogen.
  - d.) None of the above.
15. Which of the following units of measure would you use to best describe the size of a bacterium?
- a.) Meters
  - b.) Inches
  - c.) Nanometers
  - d.) Centimeters
16. As a doctor, you have two patients with the same symptoms (cough, fever, and difficulty breathing). After taking a sample of their blood, you realize they have the same illness but one has been infected with a bacteria and the other with a virus. These patients are most likely infected with which of the following illnesses?
- a.) Ebola
  - b.) Pneumonia
  - c.) E-Coli
  - d.) Smallpox

### **Appendix D: Adapted Perceived Interest Questionnaire**

In this survey we want you to rate how you responded to CRYSTAL ISLAND overall. Please indicate how strongly you agree or disagree with each statement using the f-point scale shown below.

1 Strongly Disagree      2 Disagree      3 Neutral      4 Agree      5 Strongly Agree

1. I thought CRYSTAL ISLAND was very interesting.
2. I'd like to discuss CRYSTAL ISLAND with others at some point.
3. I would play CRYSTAL ISLAND again if I had the chance.
4. I got absorbed playing CRYSTAL ISLAND without trying to.
5. I will probably think about what I learned playing CRYSTAL ISLAND for some time to come.
6. I thought CRYSTAL ISLAND's topic was fascinating.
7. CRYSTAL ISLAND was personally relevant to me.
8. I would like to play more games like CRYSTAL ISLAND in the future.
9. CRYSTAL ISLAND was one of the more interesting games I have played in a long time.
10. CRYSTAL ISLAND really grabbed my attention.

## Appendix E: CRYSTAL ISLAND Handout Materials

### Crystal Island Storyworld

The CRYSTAL ISLAND storyworld is situated on a recently discovered volcanic island where a research station (Figure 1) has been established to study the island's unique flora and fauna. There are eight main characters in the CRYSTAL ISLAND storyworld (see the "Crystal Island Characters" handout): Robert Campbell (bacteria specialist), Elise Johnson (lab technician), Kim Lee (camp nurse), Teresa Moore (senior scientist), Quentin Nash (cook and custodian), Alex Reid (player), Bryce Reid (lead scientist), and Ford Patterson (virus specialist). The user plays the role of Alex Reid visiting her father, Bryce Reid, who serves as the research station's lead scientist.



Figure 1: Crystal Island Research Station

The research camp includes the following buildings (see the "Crystal Island Virtual Environment Map" handout): *Bryce's Quarters*, the *Dining Hall*, the *Infirmary*, the *Laboratory*, and the *Living Quarters*. There is also a *waterfall* at one end of the camp.

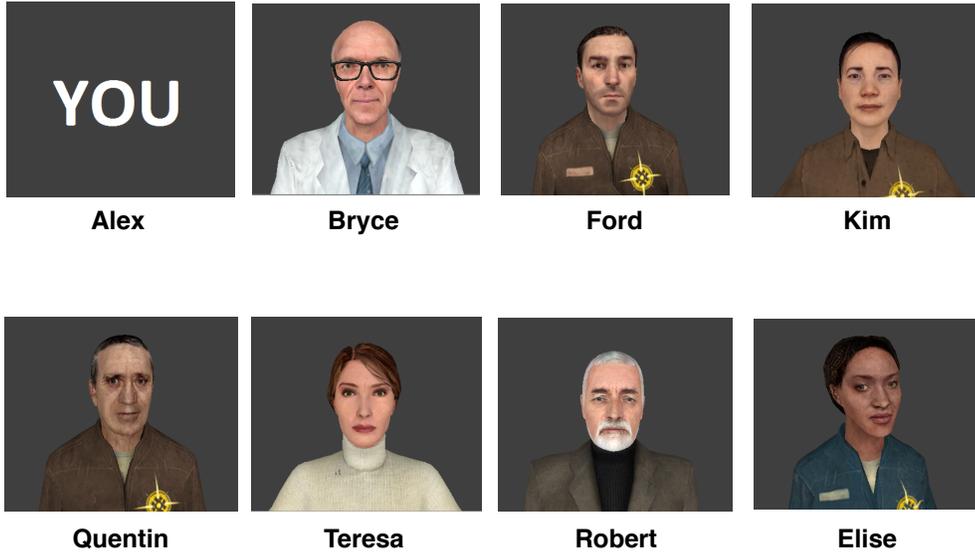
The CRYSTAL ISLAND virtual environment features a science mystery in which the user plays the role of a "medical detective". As members of the research team fall ill, it is the user's task to discover the cause of the outbreak and its source. Solve the mystery before you, and the team, run out of time!

You will be given 50 minutes to interact with the CRYSTAL ISLAND virtual environment.

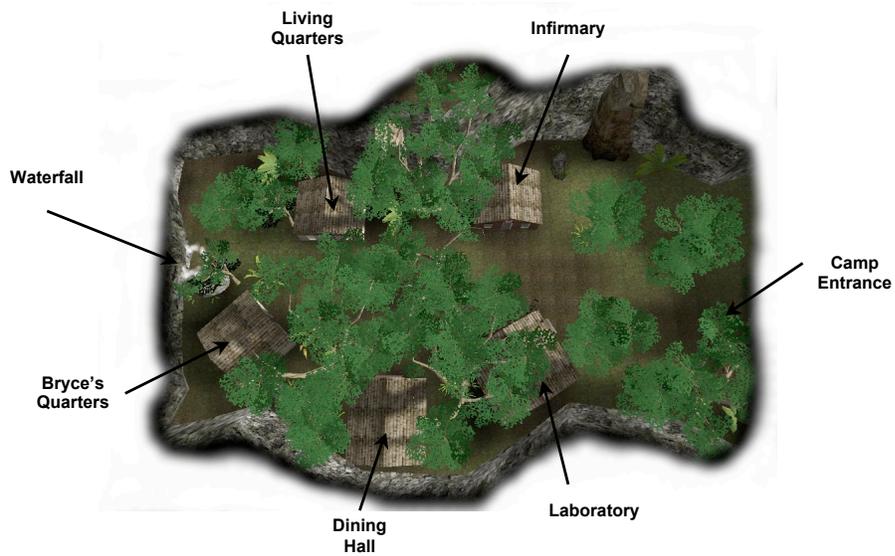
Feel free to use the "Interacting with the Virtual Environment" handout (describing the keyboard and mouse controls), as well as the "Crystal Island Virtual Environment Map" handout and the "Crystal Island Characters" handout, throughout your interaction.

If you have any questions at this time, please ask.

## Crystal Island Characters



## Crystal Island Virtual Environment Map



- |                          |                                            |                         |                                  |
|--------------------------|--------------------------------------------|-------------------------|----------------------------------|
| <b>Waterfall:</b>        | A nice place to cool off.                  | <b>Laboratory:</b>      | Research lab facilities          |
| <b>Bryce's Quarters:</b> | Living quarters for Bryce, lead scientist. | <b>Living Quarters:</b> | Living quarters for team members |
| <b>Infirmary:</b>        | Medical facilities for the camp            | <b>Camp Entrance:</b>   | Gateway to camp                  |
| <b>Dining Hall:</b>      | Camp's dining facilities                   |                         |                                  |

## Interacting with the Virtual World

### Keyboard Controls

<b>W</b> – Diagnosis worksheet		
<b>C</b> – Handheld Communicator		<b>↑</b> - Move Forward
<b>←</b> - Move Left (strafe)	<b>↓</b> - Move Backward	<b>→</b> - Move Right (strafe)

### Mouse Controls

<b>Left Mouse Button</b> – Action Button	<b>Right Mouse Button</b> – Stow/retrieve from backpack
<b>Mouse Up</b> – Look Up	
<b>Mouse Left</b> – Turn Left	<b>Mouse Right</b> – Turn Right
<b>Mouse Down</b> – Look Down	

**Appendix F: CRYSTAL ISLAND Curriculum Test v3.0**

Please answer each of the following questions to the best of your ability.

1. Which of the following statements best describes bacteria and viruses?
  - a.) Bacteria and viruses are BOTH considered alive.
  - b.) Bacteria are considered alive, but viruses are NOT considered alive.
  - c.) Viruses are considered alive, but bacteria are NOT considered alive.
  - d.) Viruses are NOT considered alive, and bacteria are NOT considered alive.
2. Which of the following sequences is in order from smallest size to largest size?
  - a.) Bacteria, viruses, human hair
  - b.) Human hair, bacteria, viruses
  - c.) Viruses, bacteria, human hair
  - d.) Viruses, human hair, bacteria
3. You place a biological agent under an electron microscope and observe that it does NOT have a nucleus. What type of agent might you be looking at?
  - a.) Bacterium
  - b.) Carcinogen
  - c.) Virus
  - d.) Either a virus or a bacterium
4. Your friend is feeling ill and goes to the doctor. The doctor gives your friend an antibiotic and as a result your friend begins to feel better soon afterwards. What infectious agent likely made your friend sick?
  - a.) Bacterium
  - b.) Carcinogen
  - c.) Virus
  - d.) Either a virus or a bacterium

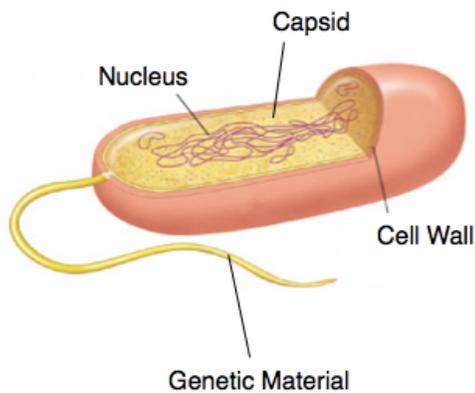
5. Your lab partners are examining a pathogen through a microscope and have observed that it is smooth and round in shape. What pathogen are your lab partners probably looking at?
  - a.) Bacterium
  - b.) Carcinogen
  - c.) Virus
  - d.) Either a virus and a bacterium
6. Which of the following statements about viruses is true?
  - a.) Viruses are considered the smallest living cells.
  - b.) Viruses consist of genetic material within a capsid.
  - c.) Viruses reproduce through binary fission.
  - d.) Virus specimens can be viewed through an optical microscope.
7. Which of the following diseases is caused by a bacterial infection?
  - a.) Ebola Hemorrhagic Fever
  - b.) Influenza
  - c.) Salmonellosis
  - d.) Smallpox
8. You are part of a team of scientists investigating an illness that may have been caused by contaminated food. Your team has asked the question, "What food source could be causing the illness?" What step should be taken next?
  - a.) Formulate a hypothesis about the source of the illness.
  - b.) Gather information about the spreading illness.
  - c.) Perform a test to identify the food source.
  - d.) Report findings about the cause of the illness.
9. What type of pathogen is considered the smallest living microorganism?
  - a.) Bacterium
  - b.) Carcinogen
  - c.) Fungi

- d.) Virus
10. An illness has been spreading through a town, and doctors have told sick people to remain at home until they are well. The spread of the illness appears to have been stopped due to the doctors' instructions. What does this tell you must be true about the illness?
- a.) It was caused by a mutagen because it spread from one organism to another.
  - b.) It was caused by a pathogen because it spread from one organism to another.
  - c.) It was caused by a carcinogen because it spread from one organism to another.
  - d.) It was not caused by a mutagen, carcinogen, or pathogen.
11. Which of the following statements about pathogens is true?
- a.) Any pathogen can be treated with antibiotics.
  - b.) Pathogens are considered living microorganisms.
  - c.) Pathogens are only responsible for a few hundred deaths each year.
  - d.) Pathogens spread from person to person.
12. You have determined that your patients' disease has been caused by a genetic mutation. Knowing this you can determine that the disease was caused by
- a.) A bacterium.
  - b.) A virus.
  - c.) A mutagen.
  - d.) None of the above.
13. Which of the following treatments is generally considered the most effective way to reduce the likelihood of a viral infection?
- a.) Antibiotic
  - b.) Chemotherapy
  - c.) Surgery
  - d.) Vaccine
14. Your friend began to feel sick this morning and is showing the following symptoms: stomach cramps, fever, and severe diarrhea. She suspects that the source of her illness

was some suspicious-looking hamburger meat she ate yesterday. Which of the following diseases is she likely suffering from?

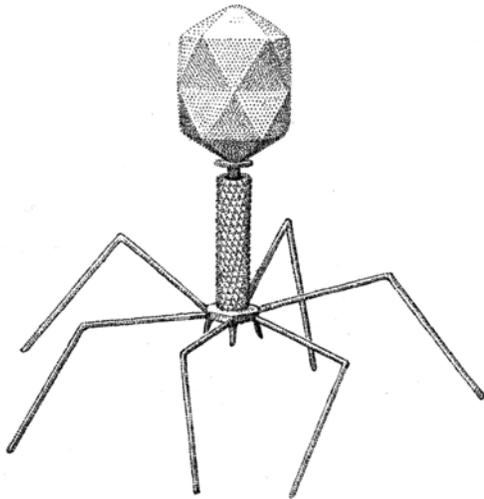
- a.) Anthrax
- b.) Botulism
- c.) Influenza
- d.) Salmonellosis

15. Bacteria come in several different shapes, but they share common structural characteristics. An illustration of one type of bacterium is shown below. Which of the illustration's labels is CORRECT?



- a.) Cell Wall
  - b.) Nucleus
  - c.) Genetic Material
  - d.) Capsid
16. Which of the following statements about viruses and bacteria is TRUE?
- a.) All viruses and bacteria are harmful to humans.
  - b.) All viruses and bacteria are considered pathogens.
  - c.) Both viruses and bacteria are composed of small cells.
  - d.) Some viruses and bacteria are not harmful to humans.
17. Which of the following characteristics do optical microscopes have in common with electron microscopes?

- a.) They both can achieve magnifications of 1,000,000x.
  - b.) They both can be used to view living specimens.
  - c.) They both can be used to view parts of bacteria specimens.
  - d.) They both use light to produce images.
18. An illustration of one type of infectious agent is shown below. Note that its shape resembles a lunar landing pod. Which type of infectious agent does the illustration most likely represent?



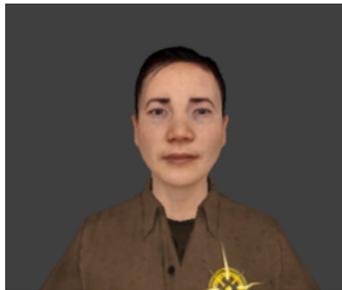
- a.) Bacterium
  - b.) Mutagen
  - c.) Virus
  - d.) Either a bacterium or a virus.
19. Which of the following diseases can cause black skin lesions?
- a.) Anthrax
  - b.) Ebola Hemorrhagic Fever
  - c.) Influenza
  - d.) Salmonellosis

**Appendix G: CRYSTAL ISLAND Narrative Questionnaire**

1. What was the main objective of your time on CRYSTAL ISLAND?
  - a.) Assist the scientific team in a dangerous research expedition.
  - b.) Identify and treat a spreading illness afflicting the research team.
  - c.) Investigate a murder that occurred on the island.
  - d.) Search for your father, a scientist who has gone missing on the island.
2. Which of the following jobs was NOT held by any team member on CRYSTAL ISLAND?
  - a.) Bacteria expert
  - b.) Cook
  - c.) Nurse
  - d.) Security guard
3. Which character initially explained the situation on CRYSTAL ISLAND, and requested your help?



a.) Elise



b.) Kim



c.) Quentin



d.) Robert

4. The condition affecting CRYSTAL ISLAND's team members came from which of the following?
  - a.) Carcinogen
  - b.) Mutagen
  - c.) Pathogen
  - d.) Physical injury
5. Which of the following symptoms did the sick team members report?
  - a.) Fever
  - b.) Muscle paralysis
  - c.) Nausea
  - d.) Rash
6. Which of the following best describes how the island's team members got sick?
  - a.) They caught an air-borne disease in the jungle.
  - b.) They consumed contaminated food.
  - c.) They were bitten by a disease-carrying insect.
  - d.) They were exposed to hazardous materials.
7. Which team member below was sick during your time on CRYSTAL ISLAND?



a.) Elise



b.) Quentin



c.) Robert



d.) Teresa

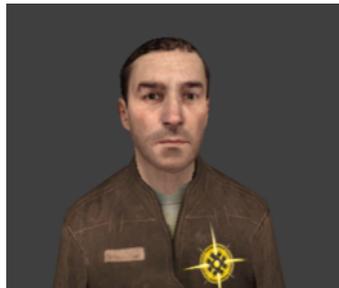
8. Which type of food had several of the sick team members consumed recently?

- a.) Apple
- b.) Coconut
- c.) Milk
- d.) Soup

9. Which of the team members had been investigating the illness before your arrival?



a. Bryce



b. Ford



c. Quentin



d. Teresa

10. After identifying a contaminated object by running tests in the laboratory, what was the best next step in CRYSTAL ISLAND's story?
- a.) Ask the lab technician to view the object under a microscope.
  - b.) Put the contaminated object in the trash can.
  - c.) Report the finding to the camp nurse.
  - d.) There were no contaminated objects on CRYSTAL ISLAND.
11. Which of the following best describes the conclusion of CRYSTAL ISLAND's story?
- a.) You administered a treatment to all of the sick team members.
  - b.) You discovered the source of the infection by running tests in the laboratory.
  - c.) You narrowly escaped CRYSTAL ISLAND before falling ill.
  - d.) You reported a diagnosis and treatment plan to the camp nurse.
12. The illness most likely affecting the team members was
- a.) Smallpox.
  - b.) Ebola hemorrhagic fever.
  - c.) Salmonellosis.
  - d.) Influenza.
13. What was the appropriate treatment for the CRYSTAL ISLAND team members?
- a.) Provide antibiotics to the sick team members.
  - b.) Provide vaccinations the healthy team members to avoid further infection.
  - c.) Recommend strict rest to the sick team members.
  - d.) Recommend chemotherapy treatments to the sick team members.

### Appendix H: Induced Narrative Adaptation Policies

The induced narrative adaptation policies for the 12 AESs in CRYSTAL ISLAND are shown in the following tables. The set of narrative-centered tutorial planner states mapped to each narrative adaptation (i.e., planner action) are shown in the columns. States that were not observed in the training data collection are omitted. For these unobserved states, the narrative-centered tutorial planner recommended a default choice of “Action 1.”

<b>Bryce's Revelation</b>		<b>Quentin's Revelation</b>	
Action 1: Deliver	Action 2: Not Deliver	Action 1: Deliver	Action 2: Not Deliver
[1,0,0,0,0,0,1,0]	[1,0,0,0,0,0,0,0]	[1,0,0,0,0,0,0,0]	[1,0,0,0,0,0,0,1]
[1,0,0,0,0,1,0,1]	[1,0,0,0,0,0,0,1]	[1,0,0,0,1,0,0,1]	[1,0,0,0,0,0,1,1]
[1,0,0,0,1,0,0,0]	[1,0,0,0,0,0,1,1]	[1,0,1,0,0,0,1,0]	[1,0,0,0,0,1,0,1]
[1,0,0,0,1,1,0,0]	[1,0,0,0,0,1,1,1]	[1,0,1,0,0,1,0,0]	[1,0,0,0,0,1,1,0]
[1,0,0,0,1,1,0,1]	[1,0,0,0,1,0,0,1]	[1,0,1,0,1,0,1,1]	[1,0,0,0,0,1,1,1]
[1,0,0,0,1,1,1,0]	[1,0,0,0,1,0,1,1]	[1,0,1,0,1,1,1,0]	[1,0,0,0,1,0,1,1]
[1,0,0,0,1,1,1,1]	[1,0,1,0,0,1,0,0]	[1,1,0,0,0,1,1,0]	[1,0,0,0,1,1,0,1]
[1,0,1,0,1,0,1,1]	[1,0,1,0,0,1,1,1]	[1,1,0,0,1,0,0,0]	[1,0,0,0,1,1,1,0]
[1,0,1,0,1,1,0,1]	[1,0,1,0,1,1,0,0]	[1,1,0,0,1,0,1,1]	[1,0,0,0,1,1,1,1]
[1,1,0,0,0,0,1,0]	[1,1,0,0,0,0,1,1]	[1,1,1,0,0,1,1,1]	[1,0,1,0,0,1,0,1]
[1,1,0,0,0,1,1,0]	[1,1,0,0,0,1,1,1]	[1,1,1,0,1,0,1,1]	[1,0,1,0,0,1,1,1]
[1,1,0,0,1,0,0,0]	[1,1,0,0,1,0,0,1]	[1,1,1,0,1,1,0,0]	[1,0,1,0,1,1,0,1]
[1,1,0,0,1,0,1,0]	[1,1,0,0,1,1,0,0]		[1,0,1,0,1,1,1,1]
[1,1,0,0,1,0,1,1]	[1,1,0,0,1,1,0,1]		[1,1,0,0,0,0,1,0]
[1,1,1,0,0,0,0,0]	[1,1,0,0,1,1,1,0]		[1,1,0,0,0,0,1,1]
[1,1,1,0,0,1,1,1]	[1,1,0,0,1,1,1,1]		[1,1,0,0,0,1,1,1]
[1,1,1,0,1,1,0,1]	[1,1,1,0,0,0,1,0]		[1,1,0,0,1,0,1,0]
	[1,1,1,0,0,1,0,1]		[1,1,0,0,1,1,0,0]
	[1,1,1,0,0,1,1,0]		[1,1,0,0,1,1,1,0]
	[1,1,1,0,1,1,1,0]		[1,1,0,0,1,1,1,1]
	[1,1,1,0,1,1,1,1]		[1,1,1,0,0,0,1,1]
			[1,1,1,0,0,1,1,0]
			[1,1,1,0,1,1,0,1]
			[1,1,1,0,1,1,1,1]

Mystery Solution					
Action 1: Influenza & Egg	Action 2: Influenza & Milk	Action 3: Influenza & Sandwich	Action 4: Salmonellosis & Egg	Action 5: Salmonellosis & Milk	Action 6: Salmonellosis & Sandwich
-	-	[0,0,0,0,1,1,1]	[0,0,0,0,0,1,1] [0,0,0,1,0,0,1] [0,0,0,0,1,0,1,1] [0,0,0,0,1,1,1,1]	-	[0,0,0,0,1,1,0]

### Diagnosis Worksheet Feedback

Action 1: Minimal Detail	Action 2: Moderate Detail	Action 3: Maximum Detail
[1,0,1,0,0,0,1,1]	[1,0,1,0,0,0,0,1]	[1,0,1,0,0,0,0,0]
[1,0,1,0,1,0,0,0]	[1,0,1,0,0,0,1,0]	[1,0,1,0,0,1,0,0]
[1,0,1,0,1,0,0,1]	[1,0,1,0,0,1,0,1]	[1,0,1,0,0,1,1,0]
[1,0,1,0,1,0,1,1]	[1,0,1,0,0,1,1,1]	[1,1,1,0,0,0,0,0]
[1,0,1,0,1,1,0,0]	[1,0,1,0,1,0,1,0]	[1,1,1,0,0,0,1,0]
[1,0,1,0,1,1,1,0]	[1,0,1,0,1,1,0,1]	[1,1,1,0,0,1,1,1]
[1,1,1,0,0,0,0,1]	[1,0,1,0,1,1,1,1]	[1,1,1,0,1,0,1,0]
[1,1,1,0,0,1,0,0]	[1,1,1,0,0,0,1,1]	[1,1,1,0,1,0,1,1]
[1,1,1,0,1,0,0,0]	[1,1,1,0,0,1,0,1]	[1,1,1,0,1,1,1,1]
	[1,1,1,0,0,1,1,0]	
	[1,1,1,0,1,0,0,1]	
	[1,1,1,0,1,1,0,0]	
	[1,1,1,0,1,1,0,1]	
	[1,1,1,0,1,1,1,0]	

### Initial Lab Test Count

Action 1: 3 Tests	Action 2: 5 Tests	Action 3: 10 Tests
[1,0,0,0,0,0,1,1]	[0,0,0,0,0,0,1,1]	[0,0,0,0,1,1,1,1]
[1,0,1,0,0,0,0,1]	[0,0,0,0,1,0,1,1]	[1,0,0,0,0,0,0,1]
[1,0,1,0,1,1,0,1]	[1,0,0,0,0,1,0,1]	[1,0,0,0,0,1,1,1]
	[1,0,0,0,1,1,0,1]	[1,0,0,0,1,0,0,1]
	[1,0,1,0,1,0,1,1]	[1,0,0,0,1,0,1,1]
	[1,0,1,0,1,1,1,1]	[1,0,0,0,1,1,1,1]
		[1,0,1,0,0,1,0,1]
		[1,0,1,0,0,1,1,1]
		[1,0,1,0,1,0,0,1]

<b>Knowledge Quiz</b>		<b>Increase Urgency</b>	
Action 1: Receive	Action 2: Not Receive	Action 1: Deliver	Action 2: Not Deliver
[1,0,0,0,0,0,1]	[1,0,0,0,1,0,0,1]	[1,0,0,0,0,1,0,0]	[0,0,0,0,0,0,1,1]
[1,0,0,0,0,0,1,1]	[1,0,0,0,1,0,1,0]	[1,0,0,0,0,1,0,1]	[0,0,0,0,1,0,1,1]
[1,0,0,0,0,1,0,1]	[1,0,0,0,1,1,0,1]	[1,0,0,0,0,1,1,0]	[1,0,0,0,0,0,0,0]
[1,0,0,0,0,1,1,0]	[1,0,0,0,1,1,1,1]	[1,0,0,0,0,1,1,1]	[1,0,0,0,0,0,0,1]
[1,0,0,0,0,1,1,1]	[1,0,1,0,0,0,0,0]	[1,0,0,0,1,0,1,1]	[1,0,0,0,0,0,1,0]
[1,0,0,0,1,0,0,0]	[1,0,1,0,0,1,1,1]	[1,0,0,0,1,1,1,0]	[1,0,0,0,0,0,1,1]
[1,0,0,0,1,0,1,1]	[1,0,1,0,1,0,0,0]	[1,0,1,0,0,0,1,0]	[1,0,0,0,1,0,0,0]
[1,0,0,0,1,1,1,0]	[1,0,1,0,1,0,1,1]	[1,0,1,0,0,1,0,1]	[1,0,0,0,1,0,0,1]
[1,0,1,0,0,0,1,0]	[1,0,1,0,1,1,0,1]	[1,0,1,0,1,0,1,1]	[1,0,0,0,1,0,1,0]
[1,0,1,0,0,0,1,1]	[1,1,0,0,0,0,0,0]	[1,0,1,0,1,1,0,1]	[1,0,0,0,1,1,0,0]
[1,0,1,0,0,1,0,0]	[1,1,0,0,0,0,1,0]	[1,0,1,0,1,1,1,1]	[1,0,0,0,1,1,0,1]
[1,0,1,0,0,1,0,1]	[1,1,0,0,0,1,1,0]	[1,1,0,0,0,0,0,0]	[1,0,0,0,1,1,1,1]
[1,0,1,0,1,0,0,1]	[1,1,0,0,1,1,0,0]	[1,1,0,0,0,0,0,1]	[1,0,1,0,0,0,0,0]
[1,0,1,0,1,1,1,1]	[1,1,1,0,0,1,0,0]	[1,1,0,0,1,0,0,0]	[1,0,1,0,0,0,0,1]
[1,1,0,0,0,0,0,1]	[1,1,1,0,0,1,1,0]	[1,1,0,0,1,0,0,1]	[1,0,1,0,0,0,1,1]
[1,1,0,0,0,0,1,1]	[1,1,1,0,0,1,1,1]	[1,1,0,0,1,0,1,1]	[1,0,1,0,0,1,0,0]
[1,1,0,0,0,1,1,1]	[1,1,1,0,1,0,0,0]	[1,1,1,0,0,0,0,0]	[1,0,1,0,0,1,1,0]
[1,1,0,0,1,0,0,0]	[1,1,1,0,1,0,0,1]	[1,1,1,0,0,0,0,1]	[1,0,1,0,0,1,1,1]
[1,1,0,0,1,0,0,1]	[1,1,1,0,1,0,1,1]	[1,1,1,0,0,0,1,0]	[1,0,1,0,1,0,0,0]
[1,1,0,0,1,0,1,0]	[1,1,1,0,1,1,0,0]	[1,1,1,0,0,0,1,1]	[1,0,1,0,1,0,0,1]
[1,1,0,0,1,1,1,0]	[1,1,1,0,1,1,1,1]	[1,1,1,0,0,1,1,0]	[1,0,1,0,1,1,0,0]
[1,1,0,0,1,1,1,1]		[1,1,1,0,0,1,1,1]	[1,0,1,0,1,1,1,0]
[1,1,1,0,0,0,0,1]		[1,1,1,0,1,0,0,0]	[1,1,0,0,0,0,1,0]
[1,1,1,0,0,0,1,0]		[1,1,1,0,1,0,0,1]	[1,1,0,0,0,1,0,1]
[1,1,1,0,0,0,1,1]		[1,1,1,0,1,0,1,0]	[1,1,0,0,0,1,1,0]
[1,1,1,0,1,1,0,1]		[1,1,1,0,1,0,1,1]	[1,1,0,0,0,1,1,1]
[1,1,1,0,1,1,1,0]		[1,1,1,0,1,1,0,0]	[1,1,0,0,1,0,1,0]
		[1,1,1,0,1,1,0,1]	[1,1,0,0,1,1,0,0]
		[1,1,1,0,1,1,1,1]	[1,1,0,0,1,1,0,1]
			[1,1,0,0,1,1,1,0]
			[1,1,0,0,1,1,1,1]
			[1,1,1,0,0,1,0,0]
			[1,1,1,0,0,1,0,1]
			[1,1,1,0,1,1,1,0]
			[1,1,1,1,1,1,1,0]

<b>Record Findings Reminder</b>	
Action 1: Receive	Action 2: Not Receive
[0,1,0,0,0,1,1,1]	[1,0,0,0,0,1,0,0]
[1,0,0,0,0,0,0,0]	[1,0,0,0,0,1,0,1]
[1,0,0,0,0,0,0,1]	[1,0,0,0,0,1,1,0]
[1,0,0,0,0,0,1,0]	[1,0,0,0,1,0,0,0]
[1,0,0,0,0,0,1,1]	[1,0,0,0,1,0,1,1]
[1,0,0,0,0,1,1,1]	[1,0,0,0,1,1,1,1]
[1,0,0,0,1,0,0,1]	[1,0,1,0,0,0,0,1]
[1,0,0,0,1,0,1,0]	[1,0,1,0,0,1,0,0]
[1,0,0,0,1,1,0,0]	[1,0,1,0,0,1,0,1]
[1,0,0,0,1,1,0,1]	[1,0,1,0,1,1,0,0]
[1,0,0,0,1,1,1,0]	[1,0,1,0,1,1,0,1]
[1,0,1,0,0,0,0,0]	[1,0,1,0,1,1,1,0]
[1,0,1,0,0,0,1,0]	[1,1,0,0,0,1,0,0]
[1,0,1,0,0,0,1,1]	[1,1,0,0,0,1,1,0]
[1,0,1,0,0,1,1,0]	[1,1,0,0,1,0,0,0]
[1,0,1,0,0,1,1,1]	[1,1,0,0,1,0,1,0]
[1,0,1,0,1,0,0,0]	[1,1,0,0,1,0,1,1]
[1,0,1,0,1,0,0,1]	[1,1,0,0,1,1,0,1]
[1,0,1,0,1,0,1,0]	[1,1,0,0,1,1,1,1]
[1,0,1,0,1,0,1,1]	[1,1,1,0,0,0,1,1]
[1,0,1,0,1,1,1,1]	[1,1,1,0,0,1,0,1]
[1,1,0,0,0,0,0,0]	[1,1,1,0,0,1,1,0]
[1,1,0,0,0,0,0,1]	[1,1,1,0,0,1,1,1]
[1,1,0,0,0,0,1,0]	[1,1,1,0,1,0,0,1]
[1,1,0,0,0,0,1,1]	[1,1,1,0,1,0,1,1]
[1,1,0,0,0,1,0,0]	[1,1,1,0,1,0,1,0]
[1,1,0,0,0,1,0,1]	[1,1,1,0,1,1,0,0]
[1,1,0,0,0,1,1,1]	[1,1,1,0,1,1,1,1]
[1,1,0,0,1,1,0,0]	[1,1,1,1,1,0,0,1]
[1,1,0,0,1,1,1,0]	
[1,1,1,0,0,0,0,0]	
[1,1,1,0,0,0,0,1]	
[1,1,1,0,0,0,1,0]	
[1,1,1,0,0,0,1,1]	
[1,1,1,0,1,0,0,0]	
[1,1,1,0,1,0,0,1]	
[1,1,1,0,1,1,0,0]	
[1,1,1,0,1,1,0,1]	
[1,1,1,0,1,1,1,0]	
[1,1,1,0,1,1,1,1]	

<b>Next Goal Prompt</b>	
Action 1: Receive	Action 2: Not Receive
[0,0,0,0,0,1,1,1]	[0,0,0,0,0,0,1,1]
[0,0,0,0,1,1,0,1]	[0,0,0,0,1,0,1,1]
[1,0,0,0,0,0,0,1]	[0,0,0,0,1,1,1,1]
[1,0,0,0,0,1,0,0]	[1,0,0,0,0,0,0,0]
[1,0,0,0,0,1,1,0]	[1,0,0,0,0,0,1,0]
[1,0,0,0,0,1,1,1]	[1,0,0,0,0,0,1,1]
[1,0,1,0,0,0,0,1]	[1,0,0,0,0,1,0,1]
[1,0,1,0,0,0,1,1]	[1,0,0,0,1,0,0,0]
[1,0,1,0,0,1,0,1]	[1,0,0,0,1,0,0,1]
[1,0,1,0,0,1,1,1]	[1,0,0,0,1,0,1,0]
[1,0,1,0,1,0,0,0]	[1,0,0,0,1,0,1,1]
[1,0,1,0,1,1,0,1]	[1,0,0,0,1,1,0,0]
[1,0,1,0,1,1,1,0]	[1,0,0,0,1,1,0,1]
[1,0,1,0,1,1,1,1]	[1,0,0,0,1,1,1,0]
[1,1,0,0,0,0,0,0]	[1,0,0,0,1,1,1,1]
[1,1,0,0,0,0,0,1]	[1,0,1,0,0,0,0,0]
[1,1,0,0,0,1,0,0]	[1,0,1,0,0,0,1,0]
[1,1,0,0,0,1,0,1]	[1,0,1,0,0,1,0,0]
[1,1,0,0,0,1,1,0]	[1,0,1,0,0,1,0,1]
[1,1,0,0,0,1,1,1]	[1,0,1,0,0,1,1,0]
[1,1,0,0,1,0,0,1]	[1,0,1,0,0,1,1,1]
[1,1,0,0,1,1,0,1]	[1,1,0,0,0,0,1,0]
[1,1,0,0,1,1,1,0]	[1,1,0,0,0,0,1,1]
[1,1,0,0,1,1,1,1]	[1,1,0,0,0,1,1,1]
[1,1,1,0,0,0,1,1]	[1,1,0,0,1,0,0,0]
[1,1,1,0,0,0,1,0]	[1,1,0,0,1,0,0,1]
[1,1,1,0,0,1,0,0]	[1,1,0,0,1,0,1,1]
[1,1,1,0,0,1,0,1]	[1,1,0,0,1,1,0,0]
[1,1,1,0,0,1,1,0]	[1,1,0,0,1,1,0,1]
[1,1,1,0,0,1,1,1]	[1,1,0,0,1,1,1,0]
[1,1,1,0,1,0,0,0]	[1,1,0,0,1,1,1,1]
[1,1,1,0,1,0,0,1]	
[1,1,1,0,1,1,0,0]	
[1,1,1,0,1,1,0,1]	
[1,1,1,0,1,1,1,0]	
[1,1,1,0,1,1,1,1]	

<b>Details of Teresa's Symptoms</b>			<b>Details of Bryce's Symptoms</b>	
Action 1: Maximum Detail	Action 2: Moderate Detail	Action 3: Minimal Detail	Action 1: Moderate Detail	Action 2: Minimal Detail
[0,0,0,0,0,0,1,1]	[0,0,0,0,1,1,1,1]	[1,0,0,0,0,0,1,0]	[0,0,0,0,1,0,1,1]	[0,0,0,0,0,0,1,1]
[0,0,0,0,0,1,1,1]	[1,0,0,0,0,0,1,1]	[1,0,0,0,0,1,1,0]	[1,0,0,0,0,0,0,1]	[0,0,0,0,0,1,1,1]
[0,0,0,0,1,0,1,1]	[1,0,0,0,1,0,0,1]	[1,0,0,0,0,1,1,1]	[1,0,0,0,0,1,0,1]	[0,0,0,0,1,1,1,0]
[1,0,0,0,0,0,0,0]	[1,0,0,0,1,0,1,1]	[1,0,0,0,1,0,0,0]	[1,0,0,0,1,0,0,1]	[0,0,0,0,1,1,1,1]
[1,0,0,0,0,1,0,1]	[1,0,0,0,1,1,1,0]	[1,0,0,0,1,0,1,0]	[1,0,0,0,1,0,1,1]	[1,0,0,0,0,0,1,1]
[1,0,0,0,1,1,0,1]	[1,0,1,0,0,1,0,1]	[1,0,0,0,1,1,0,0]	[1,0,0,0,1,1,0,0]	[1,0,0,0,0,1,1,1]
[1,0,1,0,0,0,0,0]	[1,0,1,0,1,0,0,1]	[1,0,0,0,1,1,1,1]	[1,0,0,0,1,1,0,1]	[1,0,0,0,1,0,0,0]
[1,0,1,0,0,0,0,1]	[1,0,1,0,1,0,1,1]	[1,1,0,0,0,0,1,0]	[1,0,0,0,1,1,1,0]	[1,0,1,0,0,0,1,1]
[1,0,1,0,0,0,1,1]	[1,1,0,0,1,0,1,1]	[1,1,0,0,1,0,0,0]	[1,0,0,0,1,1,1,1]	[1,0,1,0,1,0,0,1]
[1,0,1,0,0,1,1,1]	[1,1,0,0,1,1,1,0]	[1,1,0,0,1,0,0,1]	[1,0,1,0,0,1,0,0]	[1,0,1,0,1,0,1,1]
[1,0,1,0,1,1,1,1]	[1,1,1,0,0,1,1,0]	[1,1,1,0,0,1,0,0]	[1,0,1,0,0,1,0,1]	[1,0,1,0,1,1,0,0]
[1,1,0,0,0,1,1,1]	[1,1,1,0,0,1,1,1]	[1,1,1,0,1,0,1,0]	[1,0,1,0,0,1,1,1]	[1,0,1,0,1,1,1,1]
[1,1,0,0,1,1,0,0]	[1,1,1,0,1,0,0,1]	[1,1,1,0,1,1,0,0]	[1,0,1,0,1,0,0,0]	[1,1,0,0,0,0,0,1]
[1,1,0,0,1,1,1,1]	[1,1,1,0,1,1,1,0]	[1,1,1,0,1,1,1,1]	[1,1,0,0,0,0,0,0]	[1,1,0,0,0,1,1,0]
[1,1,1,0,0,0,0,0]			[1,1,0,0,0,1,1,1]	[1,1,0,0,1,0,0,0]
[1,1,1,0,0,0,1,0]			[1,1,0,0,1,0,0,1]	[1,1,0,0,1,1,0,0]
[1,1,1,0,1,0,1,1]			[1,1,0,0,1,0,1,0]	[1,1,0,0,1,1,1,0]
			[1,1,0,0,1,0,1,1]	[1,1,0,0,1,1,1,1]
			[1,1,0,0,1,1,0,1]	[1,1,1,0,0,0,1,0]
			[1,1,1,0,0,0,1,1]	[1,1,1,0,0,1,0,1]
			[1,1,1,0,0,1,1,0]	[1,1,1,0,1,0,0,0]
			[1,1,1,0,0,1,1,1]	[1,1,1,0,1,0,1,0]
			[1,1,1,0,1,0,0,1]	[1,1,1,0,1,1,0,0]
			[1,1,1,0,1,0,1,1]	[1,1,1,0,1,1,0,1]
			[1,1,1,0,1,1,1,0]	
			[1,1,1,0,1,1,1,1]	

<b>Reflection Prompt</b>	
Action 1: Deliver	Action 2: Not Deliver
[1,0,0,0,0,0,0,0]	[0,0,0,0,0,0,1,1]
[1,0,0,0,1,0,0,0]	[1,0,0,0,0,0,0,1]
[1,0,1,0,0,0,1,0]	[1,0,0,0,0,1,0,0]
[1,0,1,0,0,1,0,1]	[1,0,0,0,1,0,0,1]
[1,0,1,0,0,1,1,0]	[1,0,0,0,1,1,0,1]
[1,0,1,0,0,1,1,1]	[1,0,1,0,0,0,0,1]
[1,0,1,0,1,0,0,0]	[1,0,1,0,0,1,0,0]
[1,0,1,0,1,0,0,1]	[1,0,1,0,1,0,1,1]
[1,0,1,0,1,0,1,0]	[1,0,1,1,1,0,0,1]
[1,0,1,0,1,1,0,0]	[1,0,1,1,1,1,0,1]
[1,0,1,0,1,1,1,0]	[1,1,0,0,0,0,0,0]
[1,1,0,0,0,0,1,0]	[1,1,0,0,0,1,0,1]
[1,1,0,0,0,0,1,1]	[1,1,0,0,0,1,1,1]
[1,1,0,0,0,1,1,0]	[1,1,0,0,1,1,0,0]
[1,1,0,0,1,0,0,0]	[1,1,1,0,0,0,0,1]
[1,1,0,0,1,0,1,0]	[1,1,1,0,0,0,1,1]
[1,1,0,0,1,0,1,1]	[1,1,1,0,0,1,0,0]
[1,1,0,0,1,1,1,0]	[1,1,1,0,1,0,0,1]
[1,1,1,0,0,0,0,0]	[1,1,1,0,1,0,1,0]
[1,1,1,0,0,0,1,0]	[1,1,1,0,1,0,1,1]
[1,1,1,0,0,1,0,1]	[1,1,1,1,0,0,1,0]
[1,1,1,0,0,1,1,0]	[1,1,1,1,0,1,0,0]
[1,1,1,0,0,1,1,1]	[1,1,1,1,0,1,0,1]
[1,1,1,0,1,0,0,0]	[1,1,1,1,0,1,1,0]
[1,1,1,0,1,1,0,0]	[1,1,1,1,0,1,1,1]
[1,1,1,0,1,1,0,1]	[1,1,1,1,1,0,0,0]
[1,1,1,0,1,1,1,0]	[1,1,1,1,1,0,1,0]
[1,1,1,0,1,1,1,1]	[1,1,1,1,1,0,1,1]
[1,1,1,1,0,0,0,0]	[1,1,1,1,1,1,1,0]
[1,0,0,0,0,0,0,0]	[0,0,0,0,0,0,1,1]
[1,0,0,0,1,0,0,0]	[1,0,0,0,0,0,0,1]
[1,0,1,0,0,0,1,0]	[1,0,0,0,0,1,0,0]
[1,1,1,1,1,1,0,1]	