

# U-DIRECTOR: A Decision-Theoretic Narrative Planning Architecture for Storytelling Environments

Bradford W. Mott  
Department of Computer Science  
North Carolina State University  
Raleigh, NC 27695 USA  
bwmott@csc.ncsu.edu

James C. Lester  
Department of Computer Science  
North Carolina State University  
Raleigh, NC 27695 USA  
lester@csc.ncsu.edu

## ABSTRACT

Recent years have seen significant growth in work on interactive storytelling environments. A key challenge posed by these environments is narrative planning, in which a *director agent* orchestrates all of the events in a storyworld to create an optimal experience for a user, who is herself an active participant in the unfolding story. To create effective stories, the director agent must cope with the task's inherent uncertainty, including uncertainty about the user's intentions and the absence of a complete theory of narrative. Director agents must be efficient so they can operate in real time. In this paper, we present U-DIRECTOR, a decision-theoretic narrative planning architecture that dynamically models narrative objectives (e.g., plot progress, narrative flow), storyworld state (e.g., plot focus), and user state (e.g., goals, beliefs) with a dynamic decision network that continually selects storyworld actions to maximize *narrative utility* on an ongoing basis. The U-DIRECTOR architecture has been implemented in a narrative planner for Crystal Island, an interactive storyworld in which users play the role of a medical detective solving a science mystery. Preliminary evaluations suggest that the U-DIRECTOR architecture satisfies the real-time constraints of interactive environments and creates engaging narrative experiences.

## Categories and Subject Descriptors

H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities.

## General Terms

Algorithms, Design.

## Keywords

Interactive Narrative, Synthetic Agents.

## 1. INTRODUCTION

Narrative plays a central role in communication and cognition. With the demand for increasingly sophisticated entertainment and education applications, recent years have witnessed significant growth in research on interactive storytelling environments that create engaging narrative experiences [21, 22, 28, 33]. A critical component of this enterprise is devising effective computational models of narrative [1, 9, 20, 26, 30, 35]. These models support *director agents* that dynamically direct a cast of believable virtual characters in rich 3D storyworlds with coherent narrative structures that play out interactively in real time.

Director agents for interactive storytelling environments operate on at least two distinct but interacting levels: They craft the global story arc, typically by traversing a *plot graph* [2] that encodes a partial order of significant events in a story, and they plan the behaviors of the virtual characters and physical events in the world [3, 19]. To create an engaging experience for the user, director agents must carefully balance character believability and plot coherence [30], cope with deviations from a previously devised narrative structure caused by user actions [20, 29], provide fine-grained character control (both dialogue and actions) [22], all the while creating stories that obey the author's aesthetic, perhaps represented as an evaluation function that guides a search through a plot graph [2, 26, 35]. Director agents must also interface with believable agent functionalities supporting the story's virtual characters [14, 19, 32].

A key challenge posed by narrative planning is coping with the multiple sources of uncertainty inherent in the task. First, it is difficult to precisely infer users' goals, beliefs, and experiential attributes. It is difficult to accurately predict what effects changes in the storyworld may have on the user. Because it is important that users believe that their actions affect the ongoing story [24], predicting the effects of narrative planning actions is further complicated by the fact that users are active participants in the narrative. Second, we do not have—perhaps we *cannot* have—a complete theory of interactive narrative. We do not even have well defined theories of narrative for specific genres or particular storyworlds. Third, despite the episodic nature of narrative, however inaccurate our predictions might be over very short periods, they become even less accurate as we attempt to predict effects farther into the future.

The high degree of uncertainty, together with the multiple factors affecting narrative planning, call for a principled decision-making framework that enables director agents to rationally choose among candidate storyworld actions. The framework

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06, May 8-12, 2006, Hakodate, Hokkaido, Japan.  
Copyright 2006 ACM 1-59593-303-4/06/0005...\$5.00.

must not only be able to effectively cope with uncertainty, it must be able to exploit the evidence available about the user and storyworld state to achieve the (possibly competing) narrative objectives, and it must be able to meet the real-time performance requirements of interactive environments.

In this paper we present U-DIRECTOR,<sup>1</sup> a decision-theoretic narrative planning architecture that uses a dynamic decision network [11] to model the narrative objectives, storyworld state, and user state. Rather than adopting an *ad hoc* approach to reasoning about the multiple sources of evidence, at each tick of the clock, U-DIRECTOR systematically evaluates the available evidence, updates its beliefs, and selects the storyworld action that maximizes expected narrative utility. The U-DIRECTOR architecture has been implemented in a narrative planner for Crystal Island, an interactive storyworld in which users play the role of a medical detective solving a science mystery. Preliminary evaluations suggest that the U-DIRECTOR architecture satisfies the real-time constraints of interactive environments and creates engaging narrative experiences.

## 2. NARRATIVE PLANNING

Narrative experiences are powerful, and there is growing awareness of narrative's foundational role in psychology [7], cognitive models of reading comprehension [13], and film theory [5]. In Gerrig's seminal work on comprehension [13], he identifies two properties that readers experience. First, they are *transported*, i.e., they are somehow taken to another place and time in a manner that is so compelling it seems real. Second, they *perform* the narrative. Like actors in a play, they actively draw inferences and experience emotions as if the experiences were somehow real. Because narrative is compelling on many levels, interactive applications in entertainment, education, and training increasingly leverage narrative to create engaging experiences through rich virtual storyworld environments.

Early interactive narrative-based media were either partially or entirely scripted. All events were either completely linear or were ordered in a pre-determined branching structure [6]. However, the simplicity of the tree-like representations severely limited the level of interactivity that users could experience, and the combinatorics of highly interactive storyworlds could not be accommodated. Interestingly, the pre-scripted tradition lives on and thrives today; it is the approach used by most educational software and video games.

Dynamic narrative generation, in which stories are created on the fly (perhaps in response to users' actions), offers a promising alternative to pre-defined branching structures. Work on text-based dynamic narrative generation has yielded computational models of narrative that manipulate characters' goals and actions [23], reason about the "virtual author's" goals [18, 34], and generate extended natural language prose for narratives [8]. Dynamic narrative generation is particularly well suited to interactive storytelling environments. If narratives can be created incrementally in response to users' actions, then users can (legitimately) be made to believe that they are empowered as active participants in the story, thereby increasing their sense of agency and immersion [24].

Narrative planning for interactive storytelling environments has been the subject of increasing interest [21, 22, 28, 33].

Recent work on narrative planning has investigated a broad range of issues for interactive story environments. The narrative community has devised techniques for tightly-coupled plot creation and character behavior in dialogue-oriented interactive stories [22], search paradigms for encoding author aesthetics with an evaluation function [26, 35], and monitoring users' actions to determine if they are threatening the plot and, if so, either accommodating the new development or intervening [20, 29]. Work on emergent narratives, in which highly believable synthetic agents are given initial goals in a simulated world, has been particularly active [1, 9]. For example, a promising technique for creating emergent narratives is providing virtual characters with hierarchical task networks to drive their interactions with one another and the user [9]. Recognizing the "authoring bottleneck," some have explored efficient techniques for designing virtual characters [19, 32].

## 3. DECISION-THEORETIC NARRATIVE PLANNING

A key challenge posed by narrative planning is coping with the significant uncertainty associated with the task. First, narrative planning must deal with unobservable aspects of the user. These include her beliefs about the storyworld, her goals, and her experiential state, such as her level of engagement. Being able to effectively reason about these user characteristics is essential for proper user-story mediation. For example, it is important to have an accurate picture of the user's state to determine when and how to intervene or to accommodate the user's actions [15, 20, 29]. Second, we do not have available to us a formally represented theory of interactive narrative that supports sound and complete inference about story construction and its impact on the user. Further complicating the problems posed by uncertainty is the multitude of factors that bear on narrative decision-making activities. For example, narrative objectives such as ensuring plot progress and maintaining narrative coherence are affected by factors associated with the user's state as well as activities in the storyworld, e.g., character behaviors. Finally, narrative planning must weight all of these factors as it drives towards the goal of creating the best possible narrative experience at each juncture of the unfolding story. In short, it should satisfy the following requirement:

**Narrative Rationality:** Reasoning in a principled manner about narrative objectives, storyworld state, and user state, each with its own associated uncertainty, in the absence of a complete theory of interactive narrative, to rationally select actions that maximize expected narrative utility.

It is important to note that narrative rationality must be realized in real time to accommodate the demands of interactivity.

To address these requirements, we introduce U-DIRECTOR, a decision-theoretic *director agent* architecture that uses a dynamic decision network to achieve narrative rationality. Inspired by innovative decision-theoretic approaches to inference in intelligent tutoring systems [10, 25], U-DIRECTOR explicitly models the uncertainty in narrative objectives, storyworld state, and user state. In each decision-making cycle, it systematically evaluates the available evidence, updates its beliefs, and selects the storyworld action that maximizes expected narrative utility.

---

<sup>1</sup> U-DIRECTOR: Utility-based Director Agent

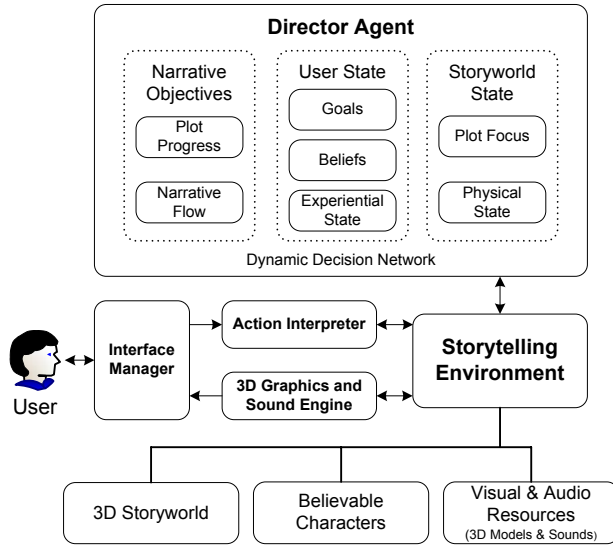


Figure 1. General Architecture.

### 3.1 Architecture

The U-DIRECTOR architecture is shown in Figure 1. All user activities in a U-DIRECTOR storytelling environment are mediated by the interface manager, which provides interaction and rendering functionalities. The storytelling environment, which drives the rendering and sound engines, employs three sets of resources: a representation of the 3D storyworld, a cast of semi-autonomous characters, and media libraries housing 3D models and sound effects. So that the director agent may carry out its primary tasks of plot creation and character behavior control, it monitors the user’s activities and the actions of the semi-autonomous characters in the storytelling environment to craft the narrative.

The director agent has access to three principle knowledge sources: narrative objectives, storyworld state, and user state. The director agent continuously monitors and seeks to achieve the *narrative objectives* to ensure that steady progress is made through the plot (plot progress) and that global coherence is maintained (narrative flow) so that the story does not seem disjoint as a result of non-motivated actions occurring. Its *storyworld state* knowledge includes information about which plot points are currently active (plot focus), as well as basic knowledge of the storyworld and characters (physical state), e.g., the characters’ goals, behaviors, and current locations. Its *user state* represents knowledge about the user’s current goals, her beliefs about the storyworld, and experiential attributes, such as engagement, which are inferred from her actions in the storytelling environment. Details of knowledge source representation are discussed in Section 3.3. To cope with the uncertainty in narrative planning, the three sets of knowledge sources are integrated into a dynamic decision network (DDN) maintained by the director agent. The director agent evaluates the DDN to solve the decision problem in each cycle to select the next narrative action.

### 3.2 Dynamic Decision Networks

Dynamic decision networks (DDNs) extend decision networks, which in turn extend Bayesian networks. *Bayesian networks* [27]

are composed of chance nodes with their associated conditional probabilities and influence arcs that collectively form a directed acyclic graph. Bayesian networks provide a compact representation of the full joint probability distribution and allow inferences to be made about any attribute within the network once priors, conditional probabilities, and available evidence have been specified.

*Decision networks* [16], also known as *influence diagrams*, extend Bayesian networks to provide a mechanism for making rational decisions by combining probability and utility theory. In decision networks, in addition to chance nodes, the network contains utility and decision nodes. The decision nodes represent the choices of the decision-maker while utility nodes model the decision-maker’s preferences. In a decision cycle, a decision theoretic agent chooses the action with the maximum expected utility.

*Dynamic decision networks* [11] provide a principled approach for agents to make rational decisions in the face of uncertainty within changing environments. To cope with time varying attributes, DDNs maintain a series of time slices to represent attributes at successive moments in time. An arc connecting an attribute in a previous time slice to an attribute in a later time slice encodes an influence on the attribute’s value from the previous attribute value. Dynamic decision networks provide a useful framework for modeling beliefs about the world, associating preferences with states of the world, and making decisions.

### 3.3 Director Agent DDNs

Narrative is fundamentally a time-based phenomenon. Director agents must therefore take into account the narrative history and be able to as accurately as possible predict (1) the effects of candidate actions on the user and (2) the effects of user’s actions on possible future courses of the narrative. Therefore, in each decision cycle, U-DIRECTOR considers candidate narrative actions to project forward in time the effects of the actions being taken and their consequent effects on the user. To do so, it evaluates its narrative objectives in light of the current storyworld state and user state. Each decision cycle considers three distinct time slices (*narrative state<sub>t</sub>*, *narrative state<sub>t+1</sub>*, and *narrative state<sub>t+2</sub>*), each of which consists of interconnected sub-networks containing chance nodes in the DDN (Figure 2). The three slices represent (1) the current narrative state, (2) the narrative state after the director agent’s decision, and (3) the narrative state after the user’s next action. The DDN’s *director action* is a decision node, the DDN’s *user action* is a chance node, and *utility<sub>t+2</sub>* is a utility node in the DDN. Each time slice encodes a probabilistic representation of the director’s beliefs about the overall state of the narrative, represented with the following knowledge sources:

- *Plot Progress*: Models the storyworld’s plot graph, a representation of temporal relations that hold between storyworld events; identifies which elements in the plot graph are waiting, ready, or completed.
- *Narrative Flow*: Models thought flow (coherence of actions to support a particular goal, e.g., searching a room after being asked to do so) and location flow (coherence of actions within spatial constraints of the storyworld, e.g., discovering multiple physical clues within the same room, one after another) associated with plot point completion; thought flow and location flow are both represented by annotations on

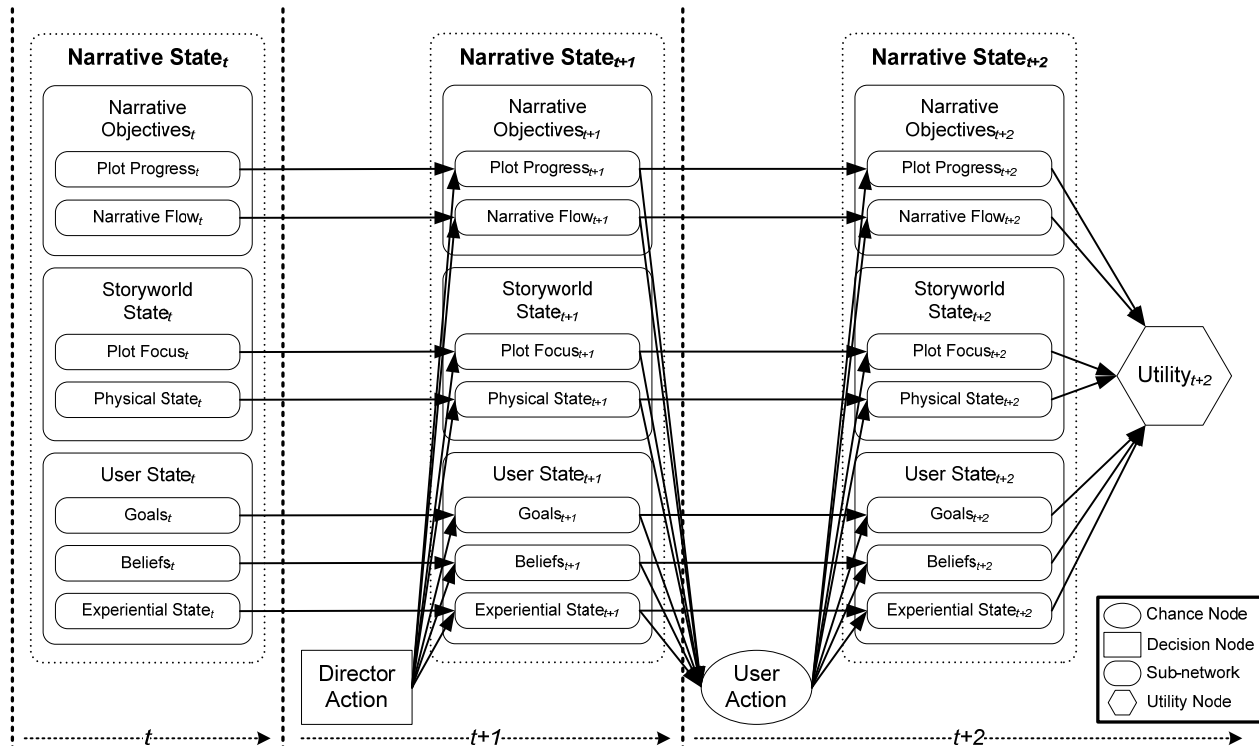


Figure 2. Director Agent Dynamic Decision Network.

user actions, indicating the relative importance for particular plot points.

- *Plot Focus*: Models the plot points that are currently active to which (it appears to the director agent) the user is currently attending.
- *Physical State*: Models the current location of the user and characters in the storyworld, and models the user’s activity level as indicated by her character’s interactions with objects and characters in the storyworld.
- *User Goals and Beliefs*: Models the user’s knowledge about the salient facts of the story that she has learned through interactions with the environment and other characters, as well as her plot progress and focus (see above).
- *User Experiential State*: Models the user’s independence (as indicated by how much manipulation the director agent has performed), her engagement (as indicated by how active she is in the environment – see physical state above), and her excitement (as indicated by changes in her knowledge about the facts of the world and pacing information).

Representations of each of these knowledge sources are integrated directly into U-DIRECTOR’s DDN. It is important to note that each knowledge source is itself in fact a sub-network encoding many beliefs about particular components of the narrative state. For example, Figure 3 depicts a portion of a plot focus network and its influences by a plot progress network. Nodes in the plot progress network are used to predict likely user actions as they make their way through the narrative. The director agent takes them into account during each decision-making cycle as described below.

Computation begins by considering the current beliefs about the narrative objectives, storyworld state, and user state, represented in *narrative state<sub>t</sub>*. U-DIRECTOR models candidate director actions and how they influence the story using links from the director action to *narrative state<sub>t+1</sub>*. Example director actions include providing various levels of hints to the user to guide them through the plot and instructing characters to perform actions that the user is either neglecting or does not seem to be capable of performing. To constrain the number of candidate actions to evaluate, U-DIRECTOR models abstract director actions [26, 35] so that it does not in this step have to attend to the plethora of concrete storyworld actions. Next, it models how the user’s actions depend on the possible worlds encoded in *narrative state<sub>t+1</sub>*. In turn, it models how possible user actions influence the story in *narrative state<sub>t+2</sub>*.

Finally, the director’s preferences over potential narrative states are modeled with links from *narrative state<sub>t+2</sub>* to the utility node *utility<sub>t+2</sub>*. Preferences provide a representation in which authors specify the relative importance of salient features of the narrative state. Narrative utility serves a similar function in the director agent’s DDN as evaluation function serve in search-based narrative planning [35]. For example, the importance of “location flow” [26] can be appropriately weighted to suit an author’s aesthetic when the DDN is constructed.

Once the director agent has fully updated the network, it selects the director action that maximizes the expected narrative utility, waits to see what action the user takes (if any) and updates its beliefs as necessary. It then begins the cycle over again, performing a “rollup” operation (also known as “filtering”), which usually involves the use of approximation techniques [4], to reduce the number of slices needed in memory at a given moment in time.

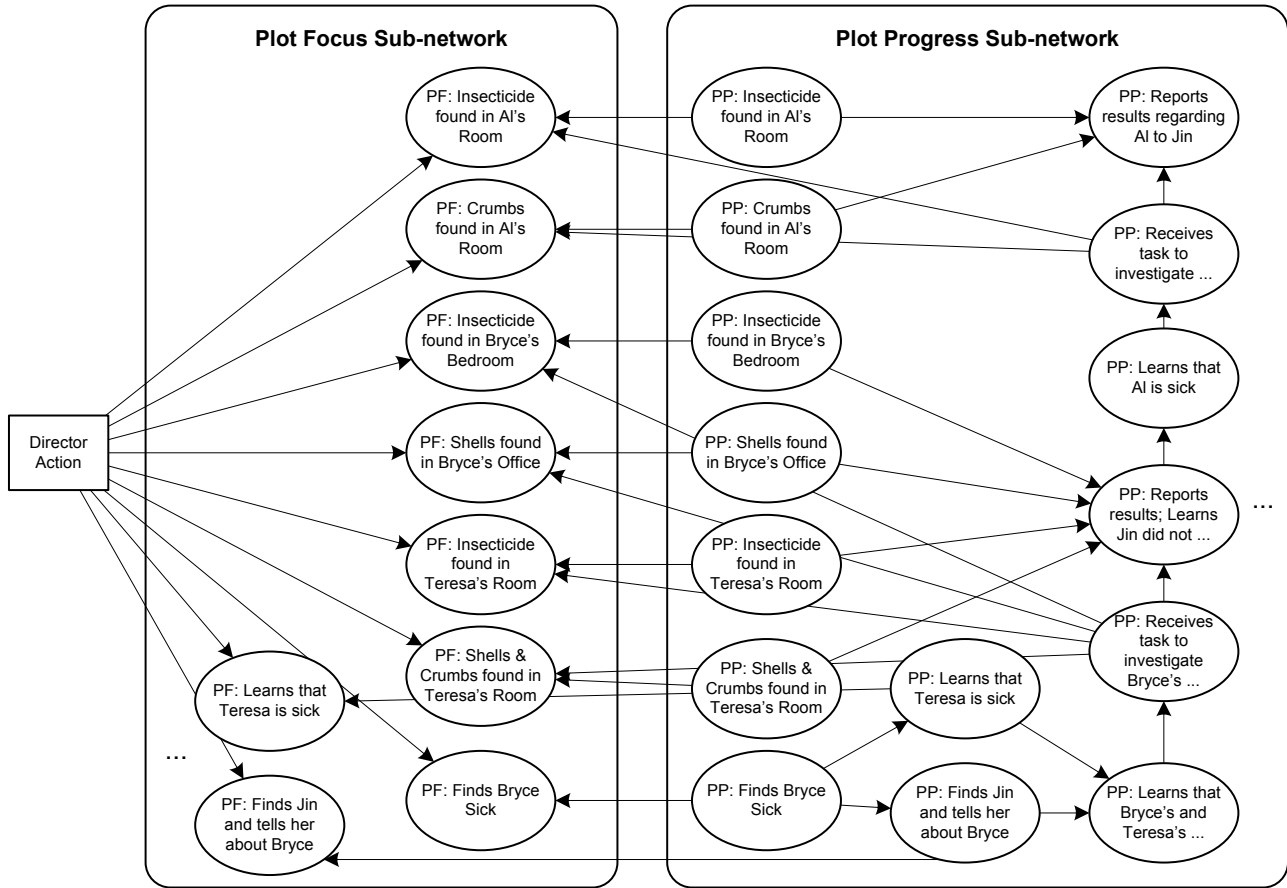


Figure 3. Selected Elements of the Plot Focus and Plot Progress Sub-networks.

## 4. EXAMPLE DOMAIN

The U-DIRECTOR architecture has been implemented and evaluated in a narrative planner for Crystal Island, an interactive storyworld in which users play the role of a medical detective solving a science mystery. After introducing the Crystal Island storyworld, we describe the implementation and present a sample interaction illustrating U-DIRECTOR’s behavior.

### 4.1 Crystal Island Storyworld

Crystal Island is a prototype interactive storytelling environment featuring a science mystery set on a recently discovered volcanic island where a research station has been established to study the unique flora and fauna. The user plays the protagonist attempting to discover the origins of an unidentified outbreak at the research station. The story opens by introducing her to the island and the members of the research team for which her father serves as the lead scientist. She is immediately confronted with her father’s sudden paralysis (Figure 4(a)). As other members of the team fall ill, it is her task to discover the cause of the outbreak as quickly as possible. She is free to explore the world to collect physical evidence and interact with other characters (Figure 4(b)). She must quickly gather enough evidence to correctly choose among candidate diagnoses including botulism, Guillain-Barré syndrome (a disorder caused by nerve inflammation), paralytic shellfish poisoning, stroke, and tick paralysis to solve the mystery and save the research group.

### 4.2 Implementation

U-DIRECTOR and the Crystal Island storytelling environment are implemented with a dynamic decision network modeled with an efficient Bayesian inference engine—a discussion of runtime performance follows in Section 5—and a high-performance 3D game platform. The director agent consists of a dynamic decision network containing approximately 200 chance nodes connected by over 400 links representing the narrative objectives, storyworld state, and user state. Collectively, the sub-networks explicitly encode over 7,000 conditional probabilities.

The director agent’s narrative utility preferences are represented in a separate structured utility network with over 50 utility nodes. The dynamic decision network for Crystal Island is modeled and implemented with the SMILE Bayesian inference engine [12], while the storytelling environment, semi-autonomous characters, and presentation layer are implemented with Valve Software’s Source™ engine, the 3D game platform for Half-Life 2.

### 4.3 Example Scenario

To illustrate U-DIRECTOR’s behavior, consider the following situation as a story unfolds on Crystal Island. The user’s character awakens from a good night’s sleep and the adventure begins. At this juncture, several elements in the plot graph are available for her to address. For example, she could find her father paralyzed in his bedroom, she might notice that a plate with

leftover food is in his office, or she might discover other facts such as that her good friend, Teresa, has also been stricken with the mysterious illness.

Although many possible actions may be taken by the user, it is the task of the director agent to determine if there is any action it might select to ensure that she is making steady progress through the plot. As the director agent assesses the situation, it notices that (1) the user is currently exhibiting high independence (since the director agent has not yet had to intervene in the story), (2) the user has not discovered any facts about the case, (3) the user is currently located in her father's house, and (4) her father is in his room. These observations, combined with a number of other factors (e.g., the fact that the user is somewhat active in the world), lead the director to make the decision to provide a hint that might lead her to find her father sick in his room since this action yields good location flow.

Once this decision is made, the system can realize the hint in many ways. For example, if the user were nearby, her father might make a moaning noise; however, if she were farther away, he might blink the lights in his room or have another character contact his daughter. The director agent then awaits the user's reaction to see how the events play out. In this case, the user hears her father's call for help, walks to his room, learns of his illness, and quickly departs to find the camp nurse. Observing that the user has found her father sick, the director agent updates its plot graph by marking the corresponding node complete and marking subsequent nodes ready for execution.

At this point the director agent must once again decide if there is anything it should do. As it notices that there are a number of plot graph elements (such as finding the partially eaten plate of food) which could be executed at the user's current location (i.e., to maintain location flow), it also takes into account the fact that these plot elements do not coincide with the user's current train of thought, as evidenced by the fact that she should be attempting to locate the camp nurse. It combines this evidence with the fact that it recently presented a hint and decides to take no action at this juncture and await the user's next action.

After some time has passed and the user has made limited progress toward completing the available plot graph elements, the director agent decides to provide a hint to the user to get her back on track. After reviewing all of the available evidence, the director agent decides that the user should be attempting to find the camp nurse and therefore directs the camp nurse to come to the user. The mystery continues with the director agent continuously monitoring the unfolding situation as the user rules out all but one of the possible diagnoses. The user solves the mystery by determining that the source of all the illnesses was botulism from last night's dinner and the research team is saved.

While the artifacts and character behaviors in the 3D world of Crystal Island and their connections to the director agent's network are the subject of continued work, the narrative reasoning described in this scenario is fully implemented. Each of the knowledge sources in the architecture (plot progress, narrative flow, plot focus, physical state, user goals, user beliefs, and user experiential state) are fully functioning and currently support this and many other scenarios.

## 5. EVALUATION

Two key aspects of narrative planning that must be addressed in a viable solution are tractability and effectiveness. Section 5.1 discusses a performance analysis of the U-DIRECTOR



(a)



(b)

Figure 4. The Crystal Island Storyworld.

implementation for Crystal Island, and Section 5.2 examines U-DIRECTOR's decision-making.

### 5.1 Performance Analysis

Real-time performance of narrative planning is critical for interactive storytelling environments. While probabilistic reasoning has much to offer, tractability in probabilistic reasoning is always a challenge. Because exact inference in Bayesian networks is known to be extraordinarily inefficient (in the worst case *NP-hard*), U-DIRECTOR exploits recent advances in approximate Bayesian inference via stochastic sampling. The accuracy of these methods depends on the number of samples used. Moreover, stochastic sampling methods typically have an "anytime" property which is particularly attractive for real-time applications.

Considering the activity level in Crystal Island and the level of abstraction at which directions are modeled, it was determined that the director agent should be able to perform a complete network update at least five times per minute. To ensure that the response time of the director agent was adequate, a performance analysis was conducted to measure the network update time using an exact Bayesian inference algorithm (*Clustering* [17]) and two

approximate Bayesian inference algorithms (*EPIS-BN* [36] and *Likelihood weighting* [31]).

Results of updating the director agent's network using the clustering, EPIS-BN, and likelihood weighting algorithms implemented in SMILE are presented in Table 1. These results were obtained on a 2.4 GHz AMD Mobile Athlon 64 PC with 1 GB of RAM running Windows XP.

**Table 1. Response times for Bayesian inference algorithms**

	Samples	Mean (seconds)	Standard Deviation
<b>Clustering</b>	N/A	18.41	8.48
<b>EPIS-BN</b>	1,000	2.32	0.18
	5,000	6.02	0.12
	10,000	10.70	0.16
<b>Likelihood</b>	1,000	1.59	0.19
	5,000	7.31	0.06
	10,000	14.65	0.03

Results of the study indicate that the clustering algorithm's running time has the greatest variability and may not satisfy the performance requirements of interactive narrative. EPIS-BN appears to provide acceptable response times while utilizing larger sample sizes than likelihood weighting and therefore yielding better approximations; however, the errors introduced by these approximation techniques can be problematic. Future work involves exploring alternative approximation techniques (e.g., [4]).

## 5.2 Simulated User Experiments

While the quality of narrative is inherently subjective, it is nonetheless important to attempt to gauge the overall effectiveness of a director agent's decision-making activities. To systematically investigate U-DIRECTOR's narrative planning, two families of simulated users were created. Simulated *cooperative* users followed the director agent's guidance 60% of the time for hints and 80% of the time for strong suggestions. In contrast, simulated *uncooperative* users followed the director agent's guidance 10% of the time for hints and 25% of the time for strong suggestions. Six simulated users (3 cooperative and 3 uncooperative) interacted with Crystal Island director agent to create six different narrative experiences.

Analyses of the traces indicate that the director agent took appropriate action to guide the user through the narrative. As was desired, the director agent tended to adopt a hint-centered approach for more cooperative users and was more heavy-handed with users who were less cooperative. The simulated user approach appears to offer a promising means for establishing baseline performance prior to conducting extensive focus group studies with human users.

## 6. CONCLUSION AND FUTURE WORK

Decision-theoretic narrative planning offers a unified approach to dynamically guiding narratives in a storytelling environment. This paper describes a decision-theoretic director agent that has been implemented in a narrative planner for an interactive storytelling environment in which the user plays the role of a medical detective. By providing a principled approach to coping with the intrinsic uncertainty of interactive narrative, decision-theoretic director agents can systematically draw inferences about the broad range of factors affecting an unfolding story. It can

thereby reason about narrative objectives, the state of the storyworld, and the state of the user to direct the course of the plot and the behaviors of the semi-autonomous characters, all with the result of creating engaging interactive narrative experiences.

Three areas of future work are particularly intriguing. First, in the current approach to constructing DDNs, devising networks by hand is very labor intensive. For example, the DDNs for the prototype interactive storytelling environment described in this paper required several person-weeks of effort. The prospect of learning the conditional probabilities in the DDNs holds much appeal.

Second, developing pedagogically oriented narrative environments that support education via narrative-based learning experiences is a promising direction for future work. However, integrating pedagogical goals into the narrative planner undoubtedly poses significant challenges because the resulting system would need to reason about (possibly competing) pedagogical and narrative goals. Nonetheless, utility-based computing could contribute to balancing these goals.

Finally, affect plays an important role in narrative. However, the current director agent includes only an impoverished representation of the user's affective state. Introducing more sophisticated techniques for affective reasoning could significantly extend the director agent's ability to accurately assess the user's engagement and motivation, which could directly contribute to its ability to create more engaging experiences.

## 7. ACKNOWLEDGMENTS

The authors wish to thank the members of the IntelliMedia Center for Intelligent Systems at NC State University for their assistance in implementing U-DIRECTOR, the Decision Systems Laboratory of the University of Pittsburgh for access to the SMILE Bayesian inference engine, Omer Sturlovich and Pavel Turzo for use of their 3D model libraries, and Valve Software for access to the Source™ engine and SDK. Thanks also to Michael Young and the members of the Liquid Narrative Group at NC State University for insightful discussions on narrative.

## 8. REFERENCES

- [1] Aylett, R. Narrative in virtual environments: towards emergent narrative. *Working notes of the Narrative Intelligence Symposium*, AAAI Spring Symposium Series. Menlo Park: AAAI Press, 1999.
- [2] Bates, J. Virtual reality, art, and entertainment. *Presence: The Journal of Teleoperators and Virtual Environments* 1(1):133-138, 1992.
- [3] Blumberg, B. and Galyean, T. Multi-level direction of autonomous creatures for real-time virtual environments. In *Proceedings of SIGGRAPH '95*, 47-54, 1995.
- [4] Boyen, X. and Koller, D. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence*, 33-42, 1998.
- [5] Branigan, E. *Narrative Comprehension and Film*. London: Routledge, 1992.
- [6] Bruckman, A. The combinatorics of storytelling: Mystery train revisited. Unpublished manuscript, 1990.

- [7] Bruner, J. *Acts of Meaning*. Cambridge, MA: Harvard University Press, 1990.
- [8] Callaway, C. and Lester, J. Narrative prose generation. *Artificial Intelligence*, 139(2), 213-252, 2002.
- [9] Cavazza, M., Charles, F., and Mead, S. Interacting with virtual characters in interactive storytelling. In *Proceedings of the First International Joint Conferences on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, 2002.
- [10] Conati, C. Probabilistic assessment of user's emotions in educational games. *Journal of Applied Artificial Intelligence*, Special Issue on Merging Cognition and Affect in HCI, 16(7-8), 555-575, 2002.
- [11] Dean, T. and Kanazawa, K. A model for reasoning about persistence and causation. *Computational Intelligence*, 5(3), 142-150, 1989.
- [12] Druzdzal, M. SMILE: structural modeling, inference, and learning engine and GeNIE: a development environment for graphical decision-theoretic models. In *Proceedings of the Eleventh Conference on Innovative Applications of Artificial Intelligence*, 902-903, 1999.
- [13] Gerrig, R. *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. New Haven: Yale University Press, 1993.
- [14] Gratch, J. and Marsella, S. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4), 2004.
- [15] Harris, J. and Young, M. Proactive mediation in plan-based narrative environments. In *Proceedings of the International Conference on Intelligent Virtual Agents*, 2005.
- [16] Howard, R. and Matheson, J. Influence diagrams. In R. Howard and J. Matheson (Eds.), *Readings on the Principles and Applications of Decision Analysis*, 721-762. Menlo Park: Strategic Decisions Group, 1984.
- [17] Huang, C. and Darwiche, A. Inference in belief networks: a procedural guide. *International Journal of Approximate Reasoning*, 15(3), 225-263, 1996.
- [18] Lebowitz, M. Story-telling as planning and learning. *Poetics*, 14, 483-502, 1985.
- [19] Loyall, B., Reilly, S., Bates, J., and Weyhrauch, P. System for authoring highly interactive, personality-rich interactive characters. In *Proceedings of the 2004 ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 2004.
- [20] Magerko, B., Laird, J., Assanie, M., Kerfoot, A., and Stokes, D. AI characters and directors for interactive computer games. In *Proceedings of the 2004 Innovative Applications of Artificial Intelligence Conference*, San Jose, CA, 2004.
- [21] Marsella, S., Johnson, L., and LaBore, C. Interactive pedagogical drama for health interventions. In *Proceeding of the 11th International Conference on Artificial Intelligence in Education*, Sydney, Australia, 2003.
- [22] Mateas, M. and Stern, A. Structuring content in the façade interactive drama architecture. In *Proceedings of Artificial Intelligence and Interactive Digital Entertainment*, Marina del Rey, 2005.
- [23] Meehan, J. Tale-spin, an interactive program that writes stories. In *Proceedings of the 5th International Joint Conferences on Artificial Intelligence*, 91-98, 1977.
- [24] Murray, J. *Hamlet on the Holodeck*. Cambridge, MA: MIT Press, 1998.
- [25] Murray, R. and VanLehn, K. DT Tutor: a dynamic, decision-theoretic approach for optimal selection of tutorial actions. In *Proceedings of the 5th International Conference on Intelligent Tutoring Systems*, Springer-Verlag, 153-162, 2000.
- [26] Nelson, M. and Mateas, M. Search-based drama management in the interactive fiction anchorhead. In *Proceedings of Artificial Intelligence and Interactive Digital Entertainment*, Marina del Rey, 2005.
- [27] Pearl J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, USA. 1988.
- [28] Prada, R., Machado, I., and Paiva, A. TEATRIX: virtual environment for story creation. In *Proceedings of the 5th International Conference on Intelligent Tutoring Systems*, Springer-Verlag, 464-473, 2000.
- [29] Riedl, M., Saretto, C., and Young, M. Managing interaction between users and agents in a multiagent storytelling environment. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems*, June, 2003.
- [30] Riedl, M. and Young, M. An intent-driven planner for multi-agent story generation. In *Proceedings of the Third International Joint Conferences on Autonomous Agents and Multi-Agent Systems*, New York, 2004.
- [31] Shachter, R. and Peot, M. Simulation approaches to general probabilistic inference on belief networks. In *Proceedings of the Fifth Annual Conference on Uncertainty in Artificial Intelligence*, 221-234, 1990.
- [32] Si, M., Marsella, S., and Pynadath, D. Thespian: using multi-agent fitting to craft interactive drama. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems*, 2005.
- [33] Swartout, W., Hill, R., Gratch, J., Johnson, L., Kyriakakis, C., Labore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thieboux, M., Tuch, L., Whitney, R., and Douglas, J. Towards the holodeck: integrating graphics, sound, character and story. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001.
- [34] Turner, S. *MINSTREL: A computer model of creativity and storytelling*. Ph.D. Thesis, University of California, Los Angeles, CA, 1992.
- [35] Weyhrauch, P. *Guiding interactive ficiton*. Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA, 1997.
- [36] Yuan, C. and Druzdzal, M. An importance sampling algorithm based on evidence pre-propagation. In *Proceedings of the Nineteenth Annual Conference on Uncertainty in Artificial Intelligence*, 624-631, 2003.