

Director Agent Intervention Strategies for Interactive Narrative Environments

Seung Y. Lee, Bradford W. Mott, and James C. Lester

Department of Computer Science, North Carolina State University
Raleigh, NC 27695, USA
{sylee, bwmott, lester}@ncsu.edu

Abstract. Interactive narrative environments offer significant potential for creating engaging narrative experiences. Increasingly, applications in education, training, and entertainment are leveraging narrative to create rich interactive experiences in virtual storyworlds. A key challenge posed by these environments is building an effective model of the intervention strategies of *director agents* that craft customized story experiences for users. Identifying factors that contribute to determining when the next director agent decision should occur is critically important in optimizing narrative experiences. In this work, a dynamic Bayesian network framework was designed to model director agent intervention strategies. To create empirically informed models of director agent intervention decisions, we conducted a Wizard-of-Oz (WOZ) data collection with an interactive narrative-centered learning environment. Using the collected data, dynamic Bayesian network and naïve Bayes models were learned and compared. The performance of the resulting models was evaluated with respect to classification accuracy and produced promising results.

Keywords: Interactive Narrative, Narrative-Centered Learning Environments, Director Agent, Dynamic Bayesian Network

1 Introduction

Recent years have witnessed significant growth in research on interactive narrative environments that create engaging narrative experiences for education, training, and entertainment applications [1,2,3,4]. Utilizing the inherent structure of narrative, interactive narrative planning [5,6,7] offers significant potential for creating engaging and effective narrative experiences. A broad range of computational models of interactive narrative planning has been investigated to build coherent narrative structures while integrating users' interactions in real-time [8,9,10,11]. A common metaphor these models share is employing a *director agent* or drama manager that works behind the scenes to direct a cast of non-player characters and storyworld events for the unfolding narrative [12,13]. Through their narrative actions, director agents subtly guide users through an intended narrative experience.

Throughout an interactive narrative, director agents actively observe the unfolding storyworld events and determine when it is most appropriate to intervene with the next director agent action to perform in service of guiding users' experiences.

Through this process, director agents manage the overall story structure and plot coherence. Prior work on interactive narrative has focused on developing models that determine the next director action to take [7,9], but little work has explicitly explored intervention strategies for director agents. A promising approach to building an effective intervention strategies model for interactive narrative is adopting an empirically driven method. By utilizing a corpus of collected human interactions within a narrative environment, models of intervention strategies can be learned from data.

This paper presents a dynamic Bayesian network (DBN) approach to modeling a data-driven director agent's intervention strategies. By utilizing multiple sources of observable evidence that affect narrative decisions from a corpus, a model of director agent intervention strategies was learned. A corpus collection was conducted using a Wizard-of-Oz methodology with thirty-three participants interacting within a narrative-centered learning environment. In the corpus collection, users assumed the role of a medical detective solving a science mystery while a wizard provided director agent functionalities for the system. Throughout the corpus collection, detailed trace data was collected for all wizard decision making and all user navigation and manipulation activities within the virtual environment. Analyses reveal that using a dynamic Bayesian network for director agent intervention strategies is promising.

This paper is structured as follows. Section 2 provides background and related work on interactive narrative. Section 3 presents the dynamic Bayesian network models for director agent intervention. The CRYSTAL ISLAND Learning Environment and the extensions to it that support Wizard-of-Oz functionalities are described in Section 4, and the data collection design and procedure are presented in Section 5. Section 6 discusses the findings and associated design implications, and Section 7 offers concluding remarks and suggests directions for future work.

2 Background

Narrative experiences are both compelling and pervasive. Cognitive scientists have long recognized narrative's foundational role in problem solving [14], and narrative plays a central role in models of reading comprehension [15] and film theory [16]. Increasingly, interactive applications in entertainment [3], education [1,4,17,18], and training [2] are leveraging narrative to create engaging experiences in rich virtual storyworlds. There has also been a long tradition of using scripted sequences of activities to create interactive narrative for education and entertainment applications. In these scripted environments, events are either completely linear or ordered in a pre-determined branching structure [19], but the tree-like representations of interactive narrative limits the level of interactivity and tailoring provided by the system to individual users.

A promising alternative to pre-defined branching structures is interactive narrative planning in which stories are dynamically crafted in response to users' actions. Interactive narrative planning has been the subject of increasing interest, and recent work in the narrative technologies community has explored a wide range of issues for interactive storytelling. Techniques have been developed for tightly coupling plot

creation and character behavior in dialogue-oriented interactive stories [3], for monitoring users' actions to determine if they are threatening the plot and, if so, either accommodating the new development or intervening [8,20]. Search-based approaches utilize an evaluation function encoding author aesthetics to guide narrative planning [7,13], while reinforcement learning has been used to learn policies to guide narrative action selection [10], and Targeted Trajectory Distribution Markov Decision Processes (TTD-MDP) have been employed to enhance replayability [11,21]. Other narrative planners have taken decision-theoretic approaches [9,22].

3 Director Agent Intervention Strategies Model

Interactive narrative is a time-dependent phenomenon. Director agents utilize numerous storyworld observations that change over time to accurately determine the most appropriate time to intervene in the unfolding story. Dynamic Bayesian networks (DBNs) explicitly characterize models' belief state over time. DBNs provide a natural representation for describing worlds that change dynamically [23]. A DBN is a directed acyclic graph that incorporates *time slices*, where each time slice contains its own state variables. By utilizing time slices, DBNs support probabilistic inference for events that change over time.

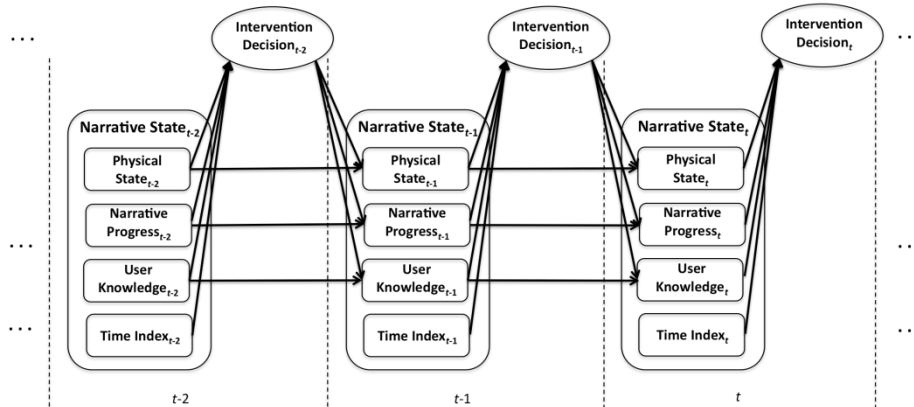


Fig. 1. Dynamic Bayesian network model for director agent intervention strategies

The high-level structure of the dynamic Bayesian network model created for director agent intervention strategies is shown in Figure 1. Three time-slices are illustrated in the figure with the intervention decision from the previous time slice, *intervention decision_{t-1}*, influencing the current intervention decision, *intervention decision_t*. Within each time slice, observations from the story world, *narrative state_t*, also influence the intervention decision. These observations include items such as the physical state of the storyworld, progression of the narrative, user knowledge of the story, and the overall story timeline. Each time slice encodes a probabilistic representation of the director agent's belief about the overall state of the narrative.

The DBN model consists of the following elements:

- *Intervention Decision*: Models the director agent intervention decision. The intervention decision is a binary variable taking on the values of either *action* or *no-action*. *Action* indicates that a director action should be taken to intervene in the story while *no-action* indicates that the director agent should remain inactive. In the network, the beliefs about *intervention decision*_{*t-1*} in time slice *t-1* are influenced by *physical state*_{*t-1*}, *narrative progress*_{*t-1*}, *user knowledge*_{*t-1*}, and *time index*_{*t-1*}. The *intervention decision*_{*t-1*} influences *physical state*_{*t*}, *narrative progress*_{*t*}, and *user knowledge*_{*t*}, which in-turn influence *intervention decision*_{*t*} in time slice *t*.
- *Physical State*: Models the current location of the characters in the storyworld's virtual environment. In the network, the current beliefs about *physical state*_{*t*} in time slice *t* are influenced by the *physical state*_{*t-1*} and *intervention decision*_{*t-1*} in time slice *t-1* and influences the *intervention decision*_{*t*} in time slice *t*.
- *Narrative Progress*: Models the storyworld's narrative structure. To characterize the progress of the narrative, we analyzed the story structure utilizing a narrative arc framework. Utilizing the current phase of the narrative arc as an observation provides the model with evidence about the high level structure of the unfolding narrative [24]. In the network, the current beliefs about *narrative progress*_{*t*} in time slice *t* are influenced by the *narrative progress*_{*t-1*} and *intervention decision*_{*t-1*} in time slice *t-1* and influences the *intervention decision*_{*t*} in time slice *t*.
- *User Knowledge*: Models the user's beliefs about the salient facts of the story learned through interactions with the environment and other characters. Within the CRYSTAL ISLAND environment users complete a *diagnosis worksheet* while solving the science mystery [25]; the diagnosis worksheet provides details regarding users' current beliefs about the story. In the network, the current beliefs about *user knowledge*_{*t*} in time slice *t* are influenced by the *user knowledge*_{*t-1*} and *intervention decision*_{*t-1*} in time slice *t-1* and influences the *intervention decision*_{*t*}.
- *Time Index*: Models the overall timeline of the storyworld to provide the temporal-based evidence for the intervention decision. In the network, *time index*_{*t*} influences the current *intervention decision*_{*t*}. It is not influenced by other observable variables since it is a deterministic monotonically increasing sequence.

During runtime, as new observations become available, such as user and wizard locations, the corresponding nodes in the network are updated with their observed values. Influences are then propagated throughout the network, allowing inferences to be made regarding the most probable *intervention decision*_{*t*} at time slice *t*.

4 Example Domain

To investigate director agent intervention strategies, a Wizard-of-Oz data collection was conducted with a customized version of the CRYSTAL ISLAND narrative-centered learning environment [24]. After introducing CRYSTAL ISLAND, we describe the custom episode created for this data collection along with the Wizard-of-Oz functionalities introduced into the environment.

4.1 CRYSTAL ISLAND

CRYSTAL ISLAND is a narrative-centered learning environment developed for middle school students for the domain of eighth-grade microbiology [4]. It is built with Valve Corporation's Source™ engine, the technology behind Half-Life® 2. CRYSTAL ISLAND features a science mystery set on a recently discovered tropical island where a research station has been established to study the island's unique flora and fauna. Within the story, the user plays the role of the protagonist attempting to discover the identity and source of an infectious disease plaguing the research station. Throughout the mystery, the user is free to explore the world and interact with other characters while forming questions, generating hypotheses, collecting data, and testing hypotheses. The user can pick up and manipulate objects, view posters, operate lab equipment, and talk with non-player characters to gather clues about the source of the disease. During the course of solving the mystery, the user completes an in-game diagnosis worksheet to organize her thoughts regarding the patients' symptoms, the likelihood of potential diseases (based on their expected symptoms, incubation period, and transmission source), and her final diagnosis. Upon completing the diagnosis worksheet, the user verifies its contents with the camp nurse and develops a treatment plan for the sickened CRYSTAL ISLAND researchers.

4.2 CRYSTAL ISLAND: Wizard-of-Oz Version

For the corpus collection, a custom episode of CRYSTAL ISLAND (Figure 2) was created featuring a companion agent who assists the user in solving the mystery. This episode features six characters Alyx Reid (player), Kim Lee (camp nurse and companion agent), Bryce Reid (lead scientist), Ford Patterson (zoologist), Audrey Newsome (botanist), Quentin Nash (camp cook), and Al Cunningham (camp foreman). The user plays the role of Alyx Reid visiting her father, Bryce, who serves as the research station's lead scientist.

Alyx has arrived at CRYSTAL ISLAND to visit her father whom she has not seen for a while. As she approaches the dock, she hears news that her father has fallen ill from Al, the camp foreman. Al tells her that Audrey, Ford, and her father were out on an expedition gathering specimens. Their expedition was scheduled to last for two days; however, they failed to return to the camp on time. Al found this very unusual since they were known to adhere closely to schedule. Fearful for their safety, Al led a search team to locate them. After two days of searching, the research team discovered that the expedition team had fallen ill on the south side of the island. It appears the

group lost their way, became ill, and could not make it back to the camp. They are in the infirmary and are being attended to by the camp's nurse. Upon hearing the news, Alyx goes to the infirmary to see her father and his colleagues. Kim, the camp's nurse, informs her that their condition is not good. Her father seems to be much worse than the others. Kim is baffled by the illness and does not know what could have caused it. She asks Alyx to help her identify the disease and its source.



Fig. 2. WOZ-enabled CRYSTAL ISLAND

In this episode, the user takes control of her character upon arriving at the camp's infirmary, which is housed in the same building as the laboratory. All of the user's interactions, as she works with the camp nurse, occur within the confines of the infirmary and laboratory. A typical scenario has the user learning about the scientific method, examining patients to learn about their symptoms, learning about infectious diseases by reading books, testing food items to find out which ones are contaminated, convincing the camp nurse of a diagnosis, and finally treating the sickened research team members for their illness.

4.3 Wizard-of-Oz Functionalities

To investigate director agent intervention strategies, CRYSTAL ISLAND was extended to include Wizard-of-Oz functionalities. In this WOZ-enabled version of CRYSTAL ISLAND, a wizard provides the narrative planning functionalities, including spoken natural language dialogue for the character of the camp nurse. Playing the role of the camp nurse, the wizard works collaboratively with the user to solve the science mystery. Together in the virtual environment they carry on rich conversations using voice chat and observe one another's actions while engaging in problem-solving activities. In addition to directing the navigation, spoken communication, and manipulation behaviors of the nurse's character in the virtual environment, the wizard guides the user's inquiry activities and controls the progression of the story. To

support these activities, the wizard's display includes detailed information regarding the user's activities in the environment (e.g., reading books, testing objects, updating the diagnosis worksheet) as well as access to a narrative dashboard. The *narrative dashboard* allows the wizard to initiate key narrative decisions in the environment (e.g., introducing new patient symptoms, having a non-player character bring in additional items for testing) analogous to a director agent.

In addition to the wizard functionalities, the narrative environment was modified to focus on the rich interactions between the user and wizard as well as to reduce the time spent navigating the environment. This was accomplished by confining the scenario to a single building housing both the camp's infirmary and laboratory. Within this environment the user and wizard gain access to all of the materials needed to solve the science mystery (e.g., sickened researchers, background books and posters, potential sources of the disease, lab equipment). The scenario, user and wizard controls, and wizard display were refined throughout a series of pilot studies with college students prior to the corpus collection described in this paper.

4.4 Example Scenario

To illustrate the behavior of the WOZ-enabled CRYSTAL ISLAND environment, consider the following scenario. A user has been collaborating with the nurse character, whose behaviors are planned and executed by the wizard. The user has learned that an infectious disease is an illness that can be transmitted from one organism to another, often through food or water. Under guidance of the nurse, the user has examined the patients' symptoms and run lab tests on food items. Through this exploration, the user has come to believe that the source of the illness is a waterborne disease and that it is likely cholera or shigellosis. Although she believes cholera is more likely, she is unable to arrive at a final diagnosis. Through her conversation with the nurse character, "Yeah, hmm, well, they both can come from water, but cholera is mostly water, I believe," the wizard determines that the user is having difficulty ruling out shigellosis and decides that this is an opportune moment to provide a hint. The wizard uses the narrative dashboard and enables the *Observe Leg Cramp Symptom* plot point, which results in one of the patients moaning loudly in the infirmary. The user examines the patient and informs the wizard, "He has leg cramps. That means it is cholera." The wizard asks the user to update her diagnosis worksheet with her new hypothesis and explain why she believes this. The user then provides a detailed explanation justifying her diagnosis, and the story concludes with the nurse treating the patients for cholera.

5 Data Collection

In the data collection, more than twenty hours of trace data were collected using the WOZ-enabled CRYSTAL ISLAND environment. The trace data includes detailed logs of all the user and wizard actions (e.g., navigation, manipulation, and decision making) within the environment, as well as audio and video recordings of their conversation.

5.1 Participants

The participants were 33 eighth-grade students (15 males and 18 females) from a public school in North Carolina ranging in age from 13 to 15 ($M = 13.79$, $SD = 0.65$). Two wizards assisted with the corpus collection, one male and one female. Each session involved a single wizard and a single student. The student and wizard were physically located in different rooms throughout the session.

5.2 Participant Procedure

When users arrived at the data collection room, they were greeted by a researcher and instructed to review a set of CRYSTAL ISLAND handouts, including information on the CRYSTAL ISLAND back-story, task description, characters, and controls. Upon completing their review of the handouts, the researcher provided further direction to the users on the use of the keyboard and mouse controls. The researcher then informed the users that they would be collaborating with another human-controlled character, the camp nurse, in the environment to solve the science mystery. Users were asked to communicate with the camp nurse throughout their sessions. Finally, the researcher answered any questions from the users, informed them that the sessions were being videotaped, instructed them to put on their headsets and position their microphones, and asked them to direct all future communication to the camp nurse. The researcher remained in the room with the user for the duration of their session. The CRYSTAL ISLAND session concluded once the user and wizard arrived at a treatment plan for the sickened researcher. The users' sessions lasted no more than sixty minutes ($M = 38$, $SD = 5.15$). During model evaluation one of the participants was eliminated as an outlier—the data were more than three standard deviations from the mean—leaving thirty-two usable trace data logs.

5.3 Wizard Protocol

To improve the consistency of the wizards' tutorial planning, narrative planning, and natural language dialogue activities, a protocol was iteratively developed and refined through a series of pilot studies. The resulting protocol included a high-level procedure for the wizard to follow (e.g., introduce yourself as the camp nurse, describe the patient situation to the student, review the scientific method with the student), a set of interaction guidelines (e.g., collaboratively work with the student to solve the mystery, organize the student's activities around the scientific method, act as a senior peer to the student, encourage the student to explain her conclusions and ensure they are logical and consistent with the available data, engage the student in constant face-to-face inquiry dialogue), and a set of narrative guidelines (e.g., overall story structure, appropriate contexts for narrative decisions, ordering constraints, be cognizant of the elapsed time to ensure the session completes in a timely manner).

Prior to the corpus collection with the eighth grade students, each wizard was trained on the CRYSTAL ISLAND microbiology curriculum and the materials that would be provided to students during the corpus collection. The wizard training also

included information on key concepts from the CRYSTAL ISLAND curriculum and the protocol to follow. After carefully reviewing the materials over the course of a week and having any of their questions answered, the wizards participated in at least three training sessions with college students. After each training session, a researcher performed an “after action review” with the wizard to discuss his or her interactions with the students and adherence to the wizard protocol.

6 Findings and Discussion

For the DBN model, there are a total of 84 time slices, 420 nodes, and more than 5200 conditional probabilities present in the director agent intervention decision-making network. The number of slices was determined based on the smallest time interval between narrative interventions found in the collected corpus. The model was implemented with the GeNIe/SMILE Bayesian modeling and inference library developed at the University of Pittsburgh’s Decision System Laboratory [26]. Given the network structure of the DBN, the probabilities of each node in the network were learned by performing parameter learning for the conditional probability tables (CPTs). The Expectation-Maximization algorithm from the SMILearn library was used to learn the CPT parameters. After CPT parameters were learned, the resulting network was used to make inferences about the director agent intervention decision nodes in the model.

6.1 Results

An analysis was conducted to assess using dynamic Bayesian networks for modeling director agent intervention strategies. To compare the effectiveness of the DBN model, a naïve Bayes model was developed as a baseline in which all observable variables are assumed to be independent of one another. Both of the models were learned using trace data collected from thirty-two interactive CRYSTAL ISLAND sessions. A leave-one-out cross validation method was employed to ensure the learned models were not over-fitted. Recall, precision, and accuracy were computed using an aggregated confusion matrix for each model.

Table 1. Classification results of director agent intervention models

Narrative Intervention Model	Recall	Precision	Accuracy
Naïve Bayes	39.3%	31.9%	75.5%
DBN	73.3%	82.0%	92.8%

Table 1 shows the results of classification measurements for the naïve Bayes and DBN models. It was found that the DBN model outperformed the baseline model naïve Bayes model in all categories. There are significant improvements exhibited by the DBN model. The DBN model achieved prediction accuracy of 92.8% and the

baseline achieved 75.5%. The DBN model exhibited a more than 16% accuracy improvement over the baseline. Also, the DBN model provides significant gains both on recall and precision, 34% and 50% respectively, as compared to the baseline.

6.2 Discussion and Design Implications

The evaluation indicates that using empirically informed dynamic Bayesian network models for director agent intervention strategies are promising. It was found that the DBN model significantly outperformed the baseline model in all classification analysis. The results suggest that in interactive narrative environments the independence assumption underlying naïve Bayes models may not hold. It appears that providing evidence regarding narrative structure, physical locations, the user's beliefs, overall story timeline, intervention decision history, and their dependent relationships can significantly improve director agent intervention decision predictions.

To further improve the prediction of the director agent intervention strategies, in addition to the observations used in the DBN model, we can utilize the *user actions* from the storyworld. The user actions indicate the user activities where they are specified by user character's interactions with objects and characters in the storyworld. User actions as an observable variable in the DBN model can influence the director agent intervention decision predictions.

It should be noted that the particular conditional probabilities learned as the parameters for the model are specific to the story arc, characters, and virtual environment of CRYSTAL ISLAND. However, it appears that the methodology itself is generalizable and applies across other narrative spaces and virtual environments. To learn the parameters for a new model, a model can be trained using the Wizard-of-Oz methodology for a new narrative space and environment, the resulting model can be used to make intervention strategies for that narrative space and virtual environment.

The results suggest design implications for modeling director agent intervention strategies. Director agents should carefully consider the state of the narrative and its emergent direction. They should also be able to reason about activities from a spatial perspective, and they must be able to utilize a model of the users' goals and beliefs so that they can properly scaffold their activities in the environment while making their intervention decisions. Furthermore, director agents should carefully leverage the conditional relationships inherit amongst the many elements of the narrative state to effectively make intervention decisions.

7 Conclusions and Future Work

Interactive narrative environments offer significant potential for crafting engaging story-based experiences that are tailored to individual users. Devising accurate models of director agent intervention strategies is critically important for creating optimal narrative experiences for users. Although previous investigations have explored techniques for modeling director agents in interactive narrative, little work has

resulted in the design of director agent intervention strategies. We have presented an empirically driven model of director agents' intervention strategies for interactive narrative-centered learning environments. A corpus collection was conducted using a Wizard-of-Oz methodology with users interacting with a WOZ-enabled version of an interactive narrative-centered learning environment. The results indicate that using empirically derived dynamic Bayesian networks for director agent intervention strategies can make accurate narrative intervention decisions.

Several directions for future work are promising. First, an important area for future work is incorporating the DBN intervention strategies model into a runtime interactive narrative system. Second, during the data collection, wizards used natural language dialogue to guide users when unexpected behaviors were encountered. Imbuing characters with sophisticated natural language dialogue capabilities offers a means for guiding users through stories, so devising adaptive models of narrative-centered interactive dialogue is a promising line of investigation. A third potentially fruitful direction is developing more effective models of affect understanding and affect generation for interactive narrative systems. Equipping narrative planners with the ability to reason about users' affective states could yield more intervention strategies that are sensitive to users' affective states that change dynamically in response to evolving narratives.

Acknowledgments. The authors wish to thank members of the IntelliMedia Group for their assistance and Valve Corporation for access to the Source™ SDK. Special thanks to Joe Grafsgaard and Kate Lester for assisting with the study. This research was supported by the National Science Foundation under Grant DRL-0822200. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M.: FearNot! – An Experiment in Emergent Narrative. In: 5th International Conference on Intelligent Virtual Agents, pp. 305–316. Kos, Greece (2005)
2. Johnson, L., Wu, S.: Assessing Aptitude for Learning with a Serious Game for Foreign Language and Culture. In: 9th International Conference on Intelligent Tutoring System, pp. 520–529. Montreal, Canada (2008)
3. Mateas, M., Stern, A.: Integrating Plot, Character, and Natural Language Processing in the Interactive Drama Façade. In: 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment, Darmstadt, Germany (2003)
4. Rowe, J., Shores, L., Mott, B., Lester, J.: Integrating Learning, Problem Solving, and Engagement in Narrative-Centered Learning Environments. *International Journal of Artificial Intelligence in Education*. In press
5. Barber, H., Kudenko, D.: Dynamic Generation of Dilemma-based Interactive Narratives. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference, Stanford, CA (2007)
6. Cavazza, M., Charles, F., Mead, S.: Interacting with Virtual Characters in Interactive Storytelling. In: 1st International Joint Conferences on Autonomous Agents and Multi-Agent Systems, pp. 318–325. Bologna, Italy (2002)

7. Nelson, M., Mateas, M.: Search-Based Drama Management in the Interactive Fiction Anchorhead. In: 1st Artificial Intelligence and Interactive Digital Entertainment Conference, Marina del Rey, CA (2005)
8. Magerko, B., Laird, J.: Building an Interactive Drama Architecture. In: 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment, pp. 226–237. Darmstadt, Germany (2003)
9. Mott, B., Lester, J.: U-Director: A Decision-Theoretic Narrative Planning Architecture for Storytelling Environments. In: 5th International Joint Conferences on Autonomous Agents and Multi-Agent Systems, pp. 997–984. Hakodate, Japan (2006)
10. Nelson, M., Roberts, D., Isbell, C., Mateas, M.: Reinforcement Learning for Declarative Optimization-Based Drama Management. In: 5th International Joint Conferences on Autonomous Agents and Multi-Agent Systems, pp. 775–782. Hakodate, Japan (2006)
11. Roberts, D., Nelson, M., Isbell, C., Mates, M., Littman, M.: Targeting Specific Distributions of Trajectories in MDPs. In: 21st National Conference on Artificial Intelligence, pp. 1213–1218. Boston, MA (2006)
12. Bates, J.: Virtual Reality, Art, and Entertainment. *Presence: The Journal of Teleoperators and Virtual Environments*. 1(1), 133–138 (1992)
13. Weyhrauch, P.: Guiding Interactive Drama. Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA (1997)
14. Bruner, J.: *Acts of Meaning*. Harvard University Press, Cambridge, MA (1990)
15. Gerrig, R.: *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. Yale University Press, New Haven (1993)
16. Branigan, E.: *Narrative Comprehension and Film*. Routledge, London (1992)
17. Marsella, S., Johnson, W.L., LaBore, C.: Interactive Pedagogical Drama for Health Interventions. In: 11th International Conference on Artificial Intelligence in Education, pp. 341–348. Sydney, Australia (2003)
18. Prada, R., Machado, I., Paiva, A.: TEATRIX: Virtual Environment for story Creation. In: 5th International Conference on Intelligent Tutoring Systems, pp. 464–473. Montreal, Canada (2000)
19. Bruckman, A.: *The Combinatorics of Storytelling: Mystery Train Revisited*. Unpublished Manuscript (1990)
20. Riedl, M., Saretto, C., Young, M.: Managing Interaction Between Users and agents in a Multi-Agent Storytelling Environment. In: 2nd International Joint Conferences on Autonomous Agents and Multi-Agent Systems, pp. 741–748. Melbourne, Australia (2003)
21. Bhat, S., Roberts, D., Nelson, M., Isbell, C., Mateas, M.: A Globally Optimal Algorithm for TTD-MDPs. In: 6th International Joint Conferences on Autonomous Agents and Multi-Agent Systems, Honolulu, Hawaii (2007)
22. Si, M., Marsella, S., Pynadath, D.: Directorial Control in a Decision-Theoretic Framework for Interactive Narrative. In: 2nd International Joint Conferences on Interactive Digital Storytelling, pp. 221–234. Guimarães, Portugal (2009)
23. Dean, T., Kanazawa, K.: A Model for Reasoning about Persistence and Causation. *Computational Intelligence*. 147(3), 142–150 (1989)
24. Lee, S., Mott, B., Lester, J.: Optimizing Story-Based Learning: An Investigation of Student Narrative Profile. In: 10th International Conference on Intelligent Tutoring System, pp. 155–165. Pittsburgh, Pennsylvania (2010)
25. Shores, L., Rowe, J., Lester, J.: Early Prediction of Cognitive Tool Use in Narrative-Centered Learning Environments. In: 15th International Conference on Artificial Intelligence in Education, pp. 320–327. Auckland, New Zealand (2011)
26. Druzzzel, M.: SMILE: Structural Modeling, Inference, and Learning Engine and Genie: A Development Environment for Graphical Decision-Theoretic Models. In: 16th National Conference on Artificial Intelligence, pp. 342–343. Orlando, Florida (1999)