

# Learning Director Agent Strategies: An Inductive Framework for Modeling Director Agents

Seung Lee, Bradford Mott, and James 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 that are tailored to individual users. 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 devising accurate models of director agents' strategies that determine the most appropriate director action to perform for crafting customized story experiences. A promising approach is developing an empirically informed model of director agents' decision-making strategies. In this paper, we propose a framework for learning models of director agent decision-making strategies by observing human-human interactions in an interactive narrative-centered learning environment. The results are encouraging and suggest that creating empirically driven models of director agent decision-making is a promising approach to interactive narrative.

## Introduction

Interactive narrative environments for education, training, and entertainment applications have been the subject of growing attention in recent years (Johnson & Wu 2008; Mateas & Stern 2005; Rowe *et al.* in press). Leveraging the inherent structure of narrative, interactive narrative planning (Barber & Kudenko 2007; Cavazza, Charles, & Mead 2002; Nelson & Mateas 2005) offers significant potential for creating engaging 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 (Magerko *et al.* 2004; Mott & Lester 2006;

Nelson *et al.* 2006; Roberts *et al.* 2006). 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 in support of the unfolding narrative (Weyhrauch 1997).

Director agents typically consider three components while crafting a global story arc for interactive narrative. They typically utilize 1) a set of *plot points* representing significant story events such as discovering a key clue while solving a mystery, 2) a set of *director agent actions* that provide a means for guiding users through an intended narrative experience, such as causing non-player characters to take action, and 3) an *evaluation function* to determine the most desirable sequence of plot points to optimize the story quality (Weyhrauch 1997; Nelson & Mateas 2005). By utilizing these components director agents strive to create engaging user experiences. Throughout an interactive narrative, director agents actively observe the unfolding story and make decisions regarding the next director agent action to perform in service of guiding the users' experience. Through this process, director agents attempt to make effective narrative decisions while managing the overall story structure and plot coherence. A significant challenge in building effective interactive narrative environments is devising accurate models of director agent's decision-making strategies (i.e., determining the next director agent action to perform and when to perform it). Prior approaches have utilized either handcrafted models of narrative decision-making (Weyhrauch 1997; Mateas & Stern 2005) or have learned models based on simulated users (Roberts *et al.* 2006).

A promising approach to building effective interactive narrative is devising empirically grounded models of director agent decision-making. Utilizing a corpus of human-human interactions collected from an interactive

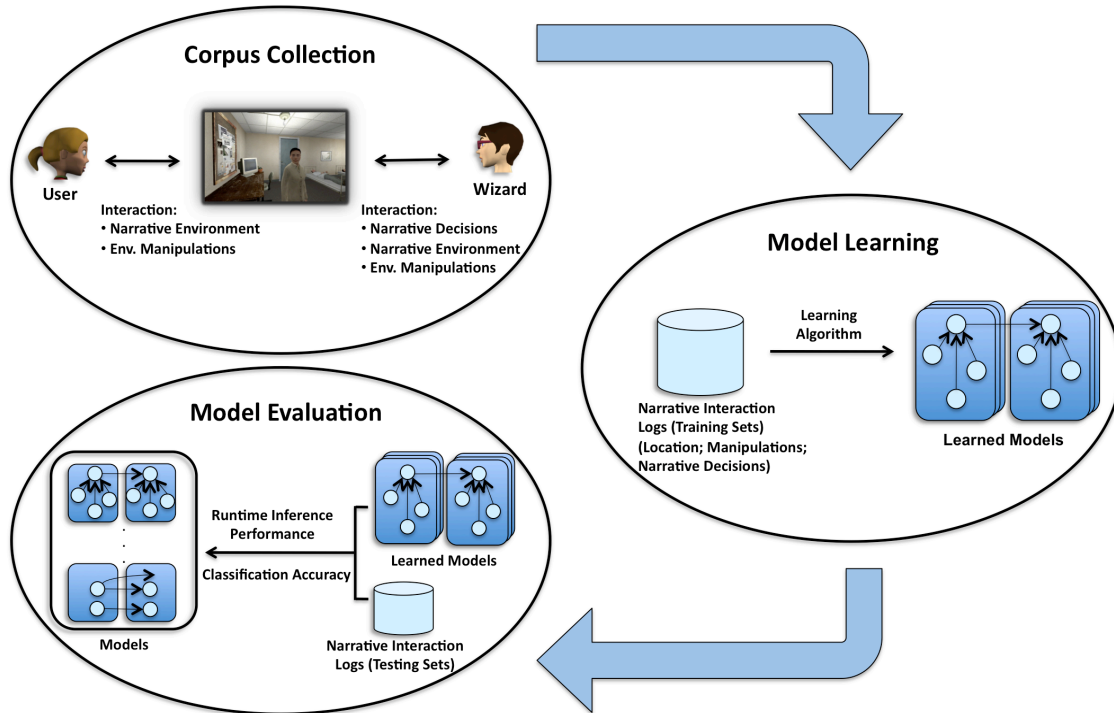


Figure 1. Data-Driven Framework for Modeling Director Agents

narrative where a human assumes the role of the director agent, models of director agent decision-making can be learned from observations of “director agents in action.”

In this paper, we present a data-driven framework for learning empirically informed director agent decision-making strategies from observations of human-human interactions within an interactive narrative environment.

## Related Work

Interactive narrative has been the subject of increasing interest, and recent work in the narrative technologies communities has explored a wide range of issues relevant to interactive storytelling environments. Techniques have been developed to create engaging experiences in rich virtual storyworlds. FAÇADE is an interactive drama that addresses the balance of character behaviors and plot creations in a dialogue oriented interactive story (Mateas & Stern 2005). In FAÇADE, the user plays the role of a guest visiting their friends, Grace and Trip, and learns that they have a serious marital problem. U-DIRECTOR is a decision-theoretic approach to narrative planning for storytelling environments (Mott & Lester 2006). By utilizing a dynamic decision network U-DIRECTOR guides actions in the virtual narrative environment. Search-Based Drama Management (SBDM) was developed by Weyhrauch (1997). SBDM is based on abstracting content of the story

into a set of plot points, choosing a set of drama manager actions, and specifying an evaluation function. Declarative Optimization-Based Drama Management (DODM) is an extension of SBDM that uses reinforcement learning to optimize its policies (Nelson *et al.* 2006). Targeted Trajectory Distribution Markov Decision Processes have been employed to enhance interactive narrative replayability (Roberts *et al.* 2006).

## Data-Driven Director Agent Framework

A data-driven framework for learning models of director agent decision-making holds much appeal. The framework consists of three-phases: corpus collection, model learning, and model evaluation (Figure 1).

### Corpus Collection

To acquire interactive narrative data that is as representative as possible, the corpus collection should address the following fundamental considerations. First, the corpus collection should be conducted using a data-collection version of the narrative environment where a wizard controls the behavior of the director agent. Beyond the wizard controlling the director agent, the user’s interaction should be as close as possible to the final version of the narrative environment. Second, to control the progression of the story, the wizard uses a *narrative*

*dashboard*, which allows the wizard to initiate director agent actions in the environment (e.g., having a non-player character perform a specific action) in the same manner that a director agent acts. In addition, the narrative dashboard provides a summary of key storyworld data that informs the wizard's decisions. It is important to note that the dashboard must be designed carefully so as not to be too complex for the wizard to effectively use while the interactive narrative is underway. Third, the wizard should make all of his or her decisions based on observable data from the environment (e.g., user behaviors in the virtual environment, dashboard data), constraints imposed on the story by the plot graph, and actions they believe lead to the most engaging and effective story along the lines of data encoded in a drama manager's evaluation function. Fourth, wizards should be given ample opportunity to become familiar with using the narrative-dashboard and the structure of the story as specified by the plot graph. Pilot data collections should be conducted to train wizards before the actual data collection

### Model Learning

In model learning, models of director agent decision-making are learned from the narrative interaction logs acquired during the corpus collection. Models are induced using a suitable machine learning algorithm based on the type of model being constructed (e.g., dynamic Bayesian network) from training datasets. To create an accurate and usable empirically informed model of director agent decision-making, a number of requirements must be satisfied. First, the model should be able to explicitly characterize the models' belief state over time and describe worlds that change dynamically since interactive narrative is a time-based phenomenon. Director agents utilize numerous storyworld observations that change over time to accurately determine their decisions. Second, the model should not only indicate the best narrative decision to make but also the appropriate time to intervene. Third, the model should be able to integrate observations from a number of sources while making decisions. Such observations might include narrative history, storyworld state, user activity, and user beliefs. Finally, the model must be capable of operating in real-time to support interactivity.

### Model Evaluation

Once the model is learned using training datasets with a machine learning algorithm, the learned model is evaluated using testing datasets to obtain the director agent's decision-making prediction performance. The performance of the learned model can be evaluated with respect to predictive accuracy. To estimate how well the learned model is performing, the model can be compared with a baseline model. The baseline model has to be carefully

selected so that proper comparisons can be made. Ultimately, once the predictive accuracy of the model is sufficient the model should be integrated into the runtime narrative environment and a user evaluation performed to ensure it operates effectively and engages users.

## Example Application

To investigate learning director agent decision-making strategies, a Wizard-of-Oz (WOZ) data collection was conducted with a customized version of the CRYSTAL ISLAND narrative-centered learning environment (Figure 2). CRYSTAL ISLAND is a virtual learning environment developed for the domain of microbiology for eighth grade science education featuring a science mystery (Rowe *et al.* in press). The WOZ-enabled CRYSTAL ISLAND extends the learning environment, using the networked multiplayer features of the Source™ Engine, to include a companion agent controlled by a wizard, who assists the user in solving the mystery.

### Modeling Approach

We propose a dynamic Bayesian network (DBN) approach to modeling director agent's decision-making strategies. The approach supports learning models from a corpus, integrating different sources of evidence affecting director agent's decisions, and the capability of operating in real-time. First, to create empirically driven models of director agent decision-making, we conducted a corpus collection with thirty-three participants using the WOZ-enabled version of CRYSTAL ISLAND. Throughout the corpus collection, detailed narrative interactions logs were collected for all wizard decision-making and all navigation and manipulation activities within the virtual environment. Second, the resulting corpus of data was utilized to learn models of director agent decision-making strategies using dynamic Bayesian networks. To create the DBN models, the network structure was manually engineered, and the parameters were learned. Each of the network structures was integrated with different narrative observable



Figure 2. WOZ-enabled CRYSTAL ISLAND

variables, and their conditional probabilities were learned using leave-one-out cross validation. The Expectation-Maximization (EM) algorithm was used to learn the conditional probability table (CPT) parameters by utilizing training datasets. The model was implemented with the GeNIe/SMILE Bayesian modeling and inference library. Finally, the learned network models were compared during an evaluation phase with a baseline model to estimate the director agent's decision-making prediction performance. The networks were analyzed for their accuracy using testing datasets.

To accurately model director agent's decision-making strategies we have developed two models: the *director intervention decision* model and the *director action decision* model. The director intervention decision model determines when is the most appropriate time for the next director agent decision to occur. The director action decision model supports making effective narrative decisions about what director agent actions to perform. Both of the models use a DBN framework with their conditional probabilities learned in the same fashion; however, the models use different storyworld observations. The learned director intervention model was compared against a naïve Bayes baseline. The DBN model achieved a prediction accuracy of 92.8% while the baseline achieved a prediction accuracy of 75.5%. The DBN model exhibited 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. The learned director action decision model was compared against a bi-gram baseline in which only the previous narrative decision was used to predict the next narrative decision. The narrative decision DBN model achieved narrative decision prediction accuracy of 93.7%. The DBN model exhibited a 23% accuracy improvement over the bi-gram model (Lee, Mott, & Lester 2011). The preliminary results indicate that a data-driven DBN model of director agent's decision-making is promising.

## Conclusions and Future Work

In this paper, we have presented a data-driven framework for an empirically grounded model of director agents' decision-making in interactive narrative environments. We outlined key requirements and considerations for corpus acquisition, model learning, and model evaluation. Finally, we summarized results of using a DBN framework for modeling director agent decision-making strategies in a narrative-centered learning environment and found that they were encouraging.

In the future, it will be important to incorporate the learned director agent decision strategies model into a

runtime interactive narrative system to validate its learned behaviors in runtime interactions with human users.

## 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

- Barber, H. and Kudenko, D. (2007). Dynamic Generation of Dilemma-based Interactive Narratives. In *Proceedings of AIIDE 2007*, 2-7. Stanford, CA.
- Cavazza, M., Charles, F., and Mead, S. 2002. Interacting with Virtual Characters in Interactive Storytelling. In *Proceedings of AAMAS 2002*, 318-325. Bologna, Italy.
- Johnson, L. and Wu, S. 2008. Assessing Aptitude for Learning with a Serious Game for Foreign Language and Culture. In *Proceedings of ITS 2008*, 520-529. Montreal, Canada.
- Lee, S., Mott, B., and Lester, J. 2011. Modeling Narrative-Centered Tutorial Decision Making in Guided Discovery Learning. In *Proceedings of AIED 2011*, 163-170. Auckland, New Zealand.
- Magerko, B., Laird, J., Assanie, M., Kerfoot, A., and Stokes, D. 2004. AI Characters and Directors for Interactive Computer Games. In *Proceedings of IAAI 2004*, San Jose, CA.
- Mateas, M. and Stern, A. 2005. Structuring Content in the Façade Interactive Drama Architecture. In *Proceedings of AIIDE 2005*, 93-98. Marina del Rey, CA.
- Mott, B., and Lester, J. 2006. U-Director: A Decision-Theoretic Narrative Planning Architecture for Storytelling Environments. In *Proceedings of AAMAS 2006*, 977-984. Hakodate, Japan.
- Nelson, M. and Mateas, M. 2005. Search-Based Drama Management in the Interactive Fiction Anchorhead. In *Proceedings of AIIDE 2005*, Marina del Rey, CA.
- Nelson, M., Roberts, D., Isbell, C., and Mateas, M. 2006. Reinforcement Learning for Declarative Optimization-Based Drama Management. In *Proceedings of AAMAS 2006*, 775-782. Hakodate, Japan.
- Roberts, D., Nelson, M., Isbell, C., Mates, M., and Littman, M. 2006. Targeting Specific Distributions of Trajectories in MDPs. In *Proceedings of AAAI 2006*, 1213-1218. Boston, MA.
- Rowe, J., Shores, L., Mott, B., and Lester, J. In press. Integrating Learning, Problem Solving, and Engagement in Narrative-Centered Learning Environments. *International Journal of Artificial Intelligence in Education*.
- Weyhrauch, P. 1997. Guiding Interactive Drama. Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA.