Computational Influence for Training and Entertainment

David L. Roberts
School of Interactive Computing
College of Computing
Georgia Institute of Technology
robertsd@cc.gatech.edu

Interactive Narrative

An *interactive narrative* is an education, training, or entertainment experience. There are two qualities that make an experience an interactive narrative. To understand those, let us clarify the meanings of interactive and narrative. Interactive: capable of acting on or influencing each other; and Narrative: presentation of events in a purposeful sequence. Thus, the qualities that define an interactive narrative are **autonomy** for players to act and **intent** of the system for the player to experience a narrative prescribed by an author.

This conception of interactive narrative is broad and encompasses many types of entertainment and training experiences. One characteristic that sets it apart from other experiences and common games like Chess is that authorial intent often requires the player's experience to be *dramatic* or adhere to some aesthetic. There is a tension between the systematic control required to ensure the intent of the narrative and the player autonomy required for interactivity—the player-driven exploration that results in an interactive quality is a potential threat to the narrative. Therefore, AI researchers in the field of interactive narrative are interested in balancing the conflicting requirements of autonomy and authorial intent. It is in this area that both I and many interactive narrative researchers have focused their efforts (Roberts and Isbell 2008). While many of the existing research projects have met with success in their own right, they have also uncovered limitations in the state of the art. I plan to complete my dissertation by addressing some of these limitations.

The tools of artificial intelligence and machine learning have often been applied in domains where the limits of human ability are stretched. For example, sophisticated AI algorithms retrieve information from the internet, filter spam from email inboxes, and aid doctors in diagnosing illnesses. In each of these cases, a technically-minded AI expert has created a technique to solve a class of problems and has put the power of that technology in the hands of a practitioner. Similarly, I plan to develop AI technologies for computer-based gaming and simulation that when given to game designers will ease their authorial burden and increase their expressive power. The fundamental question I plan to answer in my research is: How can we leverage the generalization characteristics of machine learning and ideas from social psychology to build an authoring tool that will enable domain experts to easily implement rich entertainment or training experiences through high level goal and example specification? To accomplish this, I will develop computational models of influence and implement an authoring tool for interactive experiences in order to put the power of my models in the hands of authors.

Current Work

To accomplish balance between autonomy and intent, researchers have focused on building omniscient directors, or *drama managers*, that provide a degree of authorial control while affording autonomy to players. The majority of my research to date has been focused on this area—specifically in the *Declarative Optimization-based Drama Management* (DODM) formalism.

DODM was originally proposed by Weyhrauch (1997) and later revived by Nelson and Mateas (2005). In DODM, an interactive narrative is specified with four main components: 1) a list of plot events with precedence constraints; 2) a set of abstract drama manager; 3) a model of player response to drama manager actions; and 4) an author-specified evaluation function. The drama manager's task is to select appropriate actions given the player model and the current sequence of plot events.

In its original formulation, the DODM drama manager selected actions by performing a variant of expectimax game tree search. After its revival, this approach was shown not to scale to larger narrative domains (Nelson and Mateas 2005). In order to improve the scalability of DODM, I helped develop a reinforcement learning-based approach to replace search (Nelson et al. 2006). This new technique scales significantly better but it does have a few issues of its own.

In order to achieve good results, variants of narrative domains where the drama manager was given *synthetic causers* (actions that are highly manipulative to the player) were used. As a result, the drama manager seemed to effectively cause a narrow set of narrative experiences to occur. In response to this limitation, we developed *Targeted Trajectory Distribution Markov Decision Processes* (TTD-MDPs) (Roberts et al. 2006). In TTD-MDPs, the traditional goal of optimizing some reward function is replaced with

the goal of minimizing divergence from a target distribution over experiences—the traditional notion of optimality may be sacrificed during any single episode to achieve optimality during repeated play.

Since then, there have been a number of efforts furthering the DODM formalism and TTD-MDPs. First, we developed a provably optimal algorithm for solving TTD-MDPs (Bhat et al. 2007). We have also applied the DODM framework and TTD-MDPs to guiding visitors in a museum (Roberts, Cantino, and Isbell 2007). Along slightly different lines, we have developed a technique to estimate from observations the player's satisfaction resulting from the use of a DODM-based drama manager (Roberts, Strong, and Isbell 2007). Lastly, we have developed an approach to authoring target distributions for TTD-MDPs that has very nice computational properties and we believe is more intuitive (Roberts et al. 2007).

Future Work

There have been three main criticisms of the DODM framework and TTD-MDPs: 1) DODM is abstract, divorced from concrete implementation and most of its results have been in simulation; 2) It remains to be verified that game authors can take advantage of DODM; and 3) Some feel TTD-MDPs provide variety of experience at the cost of player autonomy (Nelson and Mateas 2008). My dissertation will address all three of these concerns. In doing so, I will provide tools that enable authors of interactive narratives to create increasingly complex and engaging scenarios without increasing their burden. Progress toward this goal will enable domain experts—rather than technical experts—to leverage the power of DODM for creating interactive narratives and increase their potential to author complex and realistic content.

Computer scientists have developed powerful and widely-used algorithms by making analogies to naturally-occurring processes such as insect swarming behavior, evolution (leading to genetic algorithms/programming), hill climbing for optimization, or neural processes in artificial neural networks. Similarly, I will use the analogies of human behavior patterns identified in social psychology to develop tools for interactive narrative. I will develop a paradigm based on the ideas of influence and persuasion. By abandoning approaches based on prevention which can cause players to feel coerced and require complicated hand authored behavior models, I will enhance both the player's experience by convincing them that they are making their own decisions and the author's experience by reducing their authoring burden. To accomplish this, I will develop computational models of persuasion.

The field of social psychology has identified numerous *click-whirr* responses which function as pre-programmed behavioral responses to common social situations (Cialdini 1998). In humans, these responses appear to help us cope with the increasingly overwhelming daily stimuli that we process. A key aspect of these responses is that people believe they are making a choice by their own free will when their choice is, in effect, all but pre-determined.

By modeling persuasion, I will provide authors with powerful algorithmic tools for constructing engaging narratives. To enable non-technical authors to take advantage of these algorithms, I will build a library and example-based authoring tool that will translate high-level narrative goals into concrete implementations using click-whirr scenarios. In conclusion, by developing algorithmic models of persuasion for interactive narratives, providing a design tool to easily implement them, and evaluating that tool, I will address the criticisms of DODM by obtaining concrete results for the formalism, verifying an author's abilities to use DODM, and increasing a player's perceived autonomy.

References

Bhat, S.; Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; and Mateas, M. 2007. A globally optimal algorithm for ttd-mdps. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS07)*.

Cialdini, R. B. 1998. *Influence: The Psychology of Persuasion*. Collins.

Nelson, M. J., and Mateas, M. 2005. Search-based drama management in the interactive fiction anchorhead. In *Proc. of the 1st AI and Interactive Digital Entertainment Conf.*

Nelson, M. J., and Mateas, M. 2008. Another look at search-based drama management. In *Proc. of the 7th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*.

Nelson, M. J.; Roberts, D. L.; Isbell, C. L.; and Mateas, M. 2006. Reinforcement learning for declarative optimization-based drama management. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS06)*.

Roberts, D. L., and Isbell, C. L. 2008. A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning*. To appear.

Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; Mates, M.; and Littman, M. L. 2006. Targeting specific distributions of trajectories in mdps. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*.

Roberts, D. L.; Bhat, S.; Clair, K. S.; and Isbell, C. L. 2007. Authorial idioms for target distributions in ttd-mdps. In *Proceedings of the 22st Conference on Artificial Intelligence (AAAI-07)*.

Roberts, D. L.; Cantino, A. S.; and Isbell, C. L. 2007. Player autonomy versus designer intent: A case study of interactive tour guides. In *Proceedings of the Third Conference and Artificial Intelligence for Interactive Digital Entertainment (AIIDE-07)*.

Roberts, D. L.; Strong, C. R.; and Isbell, C. L. 2007. Estimating player satisfaction through the author's eyes. In *Proceedings of the Second Workshop on Optimizing Player Satisfaction*.

Weyhrauch, P. 1997. *Guiding Interactive Drama*. Ph.D. Dissertation, School of CS, CMU, Pittsburgh, PA. Tech Rept CMU-CS-97-109.