Opportunities for Machine Learning to Impact Interactive Narrative

David L. Roberts, Mark Riedl, and Charles L. Isbell
School of Interactive Computing, College of Computing, Georgia Institute of Technology
{robertsd,riedl,isbell}@cc.gatech.edu

Introduction

An *interactive narrative* is an education, training, or entertainment experience, based on a narrative that is guided by authorial intent. Interactive narratives have two important characteristics: some form of authorial intent as represented by the narrative structure and an affordance for autonomy of player characters. Technologies for constructing engaging interactive narratives have been springing up more and more rapidly in recent years (see [4] for a recent survey). As this field gains momentum, new and exciting opportunities arise for the application of novel machine learning methods.

In this paper, we will present a brief overview of interactive narrative, formalizing the problem in a manner amenable to computational inquiry. Then, we will survey recent work based on the use of statistical machine learning methods to solve various aspects of interactive narrative problem. We will also describe certain aspects of interactive narrative that present new and exciting challenges for the machine learning community such as non-stationarity and new definitions of optimality. Lastly, we will describe the aspects of the interactive narrative problem that we feel will both benefit the most from the application of machine learning methods as well as provide some interesting challenges for the machine learning community.

Background

There are two qualities of an entertainment or training experience that make it an interactive narrative. To understand those, let us first clarify the meanings of interactive and narrative. Interactive: capable of acting on or influencing each other; and Narrative: presentation of a series of events in a purposeful sequence, either fictional or factual. Thus, by definition, the two qualities that make an experience an interactive narrative are **autonomy** for players to act and **intent** of the system that the player have a narrative experience within bounds proscribed by the author.

This definition is broad and encompasses many types of entertainment experiences. To distinguish interactive narrative from everyday experiences and games like Chess, authorial intent often requires the player's experience to be *dramatic* or adhere to some other aesthetic. We are interested in the problem of balancing the seemingly conflicting requirements of autonomy and authorial intent. When narrative intent is threatened, the need to manage the experience in a way that balances autonomy and intent in an effective way arises. It is in this area that the majority of research in interactive narrative has been focused.

To accomplish this balance, many approaches ranging from simple `if-then` rules to modern AI and machine learning techniques such as planning [3], probabilistic graphical models [1], and reinforcement learning [2] have been used. While most of these approaches have met with success in their own right, they have also uncovered limitations in the current state of the art.

Serving Three Masters

To balance autonomy and authorial intent in an interactive narrative, one must optimize *three* distinct, possibly conflicting, dimensions: 1) player satisfaction; 2) authorial intent; 3) ease of authoring. Here, we will describe in more detail the specific challenges associated with each of these dimensions.

Player Satisfaction. While an affordance for player autonomy is one defining characteristic of interactive narrative, it does not easily lend itself to a measure of or way to specifically target player satisfaction. There are three characteristics of player preference that make its targeting difficult: 1) different players have differing preferences for play style; 2) player preferences can be (and most likely are) non-stationary—they change from episode to episode as different content is experienced; and 3) player preferences can be tacit, making their articulation difficult or impossible.

Thus, to ensure player satisfaction, there is a need to develop elicitation techniques to help construct models of player satisfaction and/or behavior. Further, to handle the changing preferences of players, a kind of "second order" model of preference change over time may need to be developed. This will help to ensure the continuing accuracy of player models during repeated episodes. Lastly, in the face of different player types with vastly different preferences, efficient and accurate recognition methods will need to be developed to ensure the proper model of preference and preference change is used during an episode.

Authorial Intent. Approaches to optimizing for the other defining characteristic, authorial intent, have included various approaches ranging from search to planning to probabilistic graphical models and beyond. In this case, the optimization is to ensure the

player's experience matches the aesthetic prescribed by the author as closely as possible. It can take on a number of forms such as "closeness" to a target or a more traditional function optimization to name a few. As intent is specified either explicitly or implicitly in the design of the narrative experience, it is often easy to provide a concrete model for a machine learning method; however, there are often subtle qualities to this intent that are, as of yet, not modeled in standard machine learning methods.

Ease of Authoring. While interactive narrative has increasingly drawn the attention of AI and ML researchers, in practice the potential authors of such experiences tend not to have much, if any, training in those disciplines. Consider asking a fiction author trained in literature to define a dynamic decision network or abstract their narrative environment into a concise STIRPS specification. This seems unlikely at best. Thus, one major challenge for the ML community is to develop techniques that are easily accessible to less technically or non-technically minded authors.

Cooperating to Not Cooperate

To further complicate things, a quality interactive experience should not rely on heavy handed intervention from the system to ensure any one of the three dimensions of desirability. As a result, the system is ideally "invisible" to the player. Thus, we have a situation contrary to existing work on modeling where it is assumed that the entity being modeled is either cooperative with the modeler or competitive against the modeler—and if not either of those, at least aware of the modeler. In an interactive narrative, the best way to achieve authorial intent and player satisfaction may be to at times create obstacles to the player's goals or to create conflict between the player and non-player characters. At other times the opposite—to reduce obstacles and conflict between characters—may be required. That is, the manager cooperates with the player to create a *globally* entertaining experience, and does so by *locally* conflicting with the player at times.

Replayability

Replayability is a thread that runs through ease of authoring, player satisfaction, and authorial intent. For player's, not repeating previous experiences helps to satisfy. For ease of authoring, avoiding repetitious authoring of alternatives and complex branching structures alleviates a significant burden. For authorial intent, authoring for multiple narrative experiences at the same time allows for players to get the most from the authors efforts. Unfortunately, machine learning methods are typically designed to optimize every single episode. In this case, the sacrifice of "optimality" during any given episode is desirable to ensure a good *distribution* over episodes is realized under repeated play.

Multi-player Influences

In the definition above, although not explicitly specified, we implied that the influence exerted as a result of interaction occurs between the player and the environment. While this may be accurate in a single player setting, when dealing with multiple players in an MMO setting, interaction occurs both between player and environment as well as between players. Thus, modeling and understanding, on a large scale, how the actions of one player can affect both the environment and the other players is crucial. Again, this is deeply tied to some of the issues pertaining to player preference described above.

Conclusion

In this paper, we have described some of the exciting opportunities for machine learning researchers to have impact in the interactive narrative community. We have highlighted some of the new challenges facing researchers. Although our discussion is brief due to space constraints, we believe we have shed light on significant avenues for future research in these areas.

References

[1] B. W. Mott and J. C. Lester. U-director: a decision-theoretic narrative planning architecture for storytelling environments. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems (AAMAS06)*, pages 977–984, New York, NY, USA, 2006. ACM Press.

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

[3] M. O. Riedl and A. Stern. Believable agents and intelligent story adaptation for interactive storytelling. In *Proceedings of the 3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE06)*, 2006.

[4] D. L. Roberts and C. L. Isbell. 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.