Player Autonomy versus Designer Intent: A Case Study of Interactive Tour Guides

David L. Roberts, Andrew S. Cantino, Charles L. Isbell
College of Computing
Georgia Institute of Technology
{robertsd,cantino,isbell}@cc.gatech.edu

Abstract

We explore the tradeoff between player autonomy and designer intent by simulating a system of autonomous museum tour guides. Visitors may have different art preferences or may wish to visit different exhibits on multiple visits. Often, these desires conflict. For example, visitors may wish to see the museum's most popular work, but that could cause congestion, ruining the experience. Thus, our task is to build a set of guides that can satisfy their visitors' goals while also providing quality experiences for all. We present the results of a case study indicating that there is a space in the design spectrum between fully author-controlled narrative and complete player autonomy that reduces the frequency of bad experiences while allowing visitors to realize their own goals.

Introduction

In this paper, we describe a system of interactive tour guides that guide visitors through a complex social environment. Our goal is to explore the tradeoff between the autonomy of human participants and the control necessary to ensure designer intent. Curators organize exhibits flow, so museums are well suited for studying designer intent. Additionally, museums are large, with complex layouts, many simultaneous visitors, opportunity for repeat visits, and a variety of visitor goals. As with other cases of interactive entertainment (IE), we are forced to balance competing desires. We want tours that enable visitors to see as much art as possible without overwhelming them. We want visitors to see art of interest, while limiting museum congestion. Finally, we want to allow visitors to exercise autonomy by sometimes ignoring the tour guide, while still ensuring that they benefit from the guide's insight. These tradeoffs are analogous to those faced by a *drama manager* that balances authorial intent and player autonomy in IE: a tour is a story (or sequence of plot events) in an interactive drama, and congestion represents undesirable game states.

Problem Description

In this section, we provide a brief overview of TTD-MDPs, the framework for decision making used by our tour guides, and then present our museum and visitor models. Additionally, we describe the experimental setup of our case study.

TTD-MDPs A Markov Decision Processes is a probabilistic state model of how agents move through an environment while making decisions. A TTD-MDP is a tuple $< \mathcal{T}, \mathcal{A}, P, P(\mathcal{T}) >$, where $\mathcal{T}$ are finite-length trajectories of states, a set of actions $\mathcal{A}$, a transition model P, and a target distribution over complete trajectories $P(\mathcal{T})$. The solution to a TTD-MDP is a policy $\pi : \mathcal{T} \rightarrow P(\mathcal{A})$ providing a distribution over actions in every state. The optimal policy results in long-term behavior as close to the target distribution as possible. The reader may find further details in (Roberts *et al.* 2006).

Modeling a Museum Our museum is a 4x5 grid of rooms, some containing interesting objects, with walls preventing some transitions. A trajectory is a sequence of rooms modeling a particular tour, where each room is of the form: (x, y, c). The c parameter indicates whether the room is over-capacity, thus "congested." We assume that guides can detect the current room and can communicate locally with nearby guides to determine surrounding room congestion.

Tour Probabilities With TTD-MDPs, the designer achieves a desired behavior by properly selecting a target probability distribution over trajectories. In the museum domain, we 1) collect a set of prototypical "good" tours—the museum curator's intent for traffic flow and visitor experience—and use them as the centroids in a Gaussian mixture model; and 2) define a distance metric, based on *edit distance* (Levenshtein 1965; 1966), to obtain target probabilities for each tour.

Modeling Visitors We model both *naive visitors* and *informed visitors*. The latter represent dedicated art spectators, with larger sets of goals. In addition, we consider two variants of these visitor types: the *new visitor* and the *returning visitor*. We model new visitors as having no history of satisfied goals, while returning visitors have 35% of the possible goals already satisfied. Note that the history of satisfied goals is only used for evaluating the quality of experience after experimentation—the tour guides are unaware of any goals that haven't been met during the current tour. In our 4x5 museum world, we select 10 of the 20 rooms to contain potential goals for the informed visitor and 6 to contain potential goals for the naive visitor. For each of the visitor types, we assign 3 goals to be "hidden" goals, or goals that the visitor will enjoy but does not know to pursue. Each in-

Measure	Congestion		New Goals		Hidden Goals	
Model	L	H	L	H	L	H
TTD:	**0.135**	**0.153**	**0.476**	**0.598**	**0.289**	**0.351**
ignore:	**0.209**	**0.202**	**0.497**	**0.608**	**0.290**	**0.374**
wander:	0.517	0.517	0.113	0.271	0.118	0.273
random:	0.287	0.247	0.398	0.554	0.226	0.342

Table 1: Aggregate statistics for low and high goal density.

stantiation of a visitor type in an experiment receives some subset of that type's potential goals as their individual goals. We consider both high and low goal density scenarios.

The tour guides lead visitors by suggesting transitions for them to make according to the distribution obtained by solving the TTD-MDP. The available suggestions are $\{north, south, east, west, no\_suggestion\}$. Visitors generally follow suggestions except when near a known goal location, when they tend to move to the goal. Further, visitors prefer not to revisit rooms. This gentle guidance model is similar to the hint model used in *Declarative Optimization-Based Drama Management* (Nelson *et al.* 2006). We divide visitors' willingness to follow suggestions into categories: those who *possibly*, *probably*, or *definitely* follow suggestions, and visitors who always *ignore* suggestions.

Experimental Design We compare results for TTD-MDP tour guides to three approaches. The first two of these approaches use no tour guide: 1) *wander*, a visitor who acts randomly, and 2) *ignore*, a visitor who pursues a goal when one step away but wanders otherwise. The third approach, *random*, uses a tour guide that chooses actions uniformly.

We require that every goal lie on at least one prototype trajectory for each visitor type. In simulation, we choose uniformly from the four visitor types and allow them to enter the museum at a constant rate of n per simulation step. Repeatedly, we allow all museum visitors to move to a neighboring room in a random order. We update room congestion states after each visitor moves.

Results

To measure the effectiveness of tour guides are at realizing visitors' goals, we consider: the percentage of a visitor's known goals and hidden goals that are achieved, and the frequency of congested rooms experienced by each visitor.

Goals and Congestion We examine both congestion and goal realization together, in aggregate for all visitor types. In Table 1, we consider the results of experiments both with and without TTD-based guides as well as with both the *wander* and *random* baselines. "L" and "H" represent low and high goal density experiments. These results are shown averaged across all visitor models (naive and informed in both the new and returning variants). The *wander* and *random* baselines do not perform well in any of the categories. In the case of the *wander* baseline, this is attributable to a lack of goal directed behavior. For the *random* tour guide, however, this is attributable to the willingness of the visitor to follow the guide's random suggestions. In comparison to those baselines, the *ignore* and TTD cases yield very promising results. Specifically, we see a noticeable reduction in congestion that accompanies a significant increase in goal realization. Note the differences between the TTD case and the ignore case (in bold). There we see that the rate of congestion is greatly reduced while goal realization is preserved when TTD-based guides are used.

Measure	Congestion		New Goals		Hidden Goals	
Capacity	L	H	L	H	L	H
inf	0	0	0.501	0.569	0.313	0.346
6	0.021	0.029	0.487	0.564	0.307	0.344
5	0.049	0.062	0.472	0.560	0.297	0.342
4	0.116	0.133	0.441	0.543	0.272	0.338
3	0.244	0.259	0.416	0.534	0.253	0.333

Table 2: Aggregate Statistics for varying room capacity limits.

Measure	Congestion		New Goals		Hidden Goals	
Model	L	H	L	H	L	H
ignore	**0.209**	**0.202**	0.497	0.608	0.290	0.374
possibly	**0.135**	**0.153**	0.476	0.598	0.289	0.351
probably	**0.116**	**0.133**	0.441	0.544	**0.274**	**0.338**
definitely	**0.091**	**0.090**	0.364	0.450	**0.315**	**0.385**

Table 3: Aggregate Statistics for Visitors with varying willingness to follow suggestions.

Capacity and Visitor Autonomy In Table 2, we summarize the effects of room capacity. In particular, we see that the effects of room capacity on goal realization are more pronounced in the low goal density case than in the high density case. As capacity decreases, the percentage of realized goals in the high density case remains essentially the same. The effect on goal realization of decreasing room capacity in the low goal density case is exaggerated because the same number of visitors are competing to achieve fewer goals. As a result of the guides' tendencies to suggest alternates to congested rooms and visitors' tendency to follow the guides, we see a reduction in goal satisfaction.

In Table 3 we present the effect of autonomy on congestion and goal realization. The data was obtained by varying the visitors' willingness to follow advice. Note that the rate of congestion is slightly lower for the low goal density case. We attribute this to visitors having fewer goals to seek, and therefore being more willing to follow tour guide advice that may lead them both away from congestion as well as away from unrealized goals. Taken together with the percentage of satisfied goals, this data is informative. We see that the more willing a visitor is to follow the tour guide, the less congestion it will encounter, but the fewer goals it will realize; however, this tradeoff may be worthwhile—a 26.0% reduction in goal satisfaction accompanies a 55.4% reduction in congestion (in the high goal density experiment).

However, although the realization frequency of hidden goals generally decreases as visitors follow their guides, if they always follow their guides they start to realize more hidden goals again. This occurs because the guides have some sense of where hidden goals may be, due to the museum curator's well constructed prototype tours.

As visitors have more autonomy, they achieve more of their goals because of their willingness to ignore the tour guide and pursue a known goal; however, this gives rise to a tragedy of the commons: when visitors always act only in their own immediate interest, they end up in crowded parts of the museum lessening the quality of the experience for everyone. On the other hand, if visitors always listen to the tour guide, they experience less congestion at the expense of realizing fewer of their known goals. Somewhere in the middle of these extremes is a “sweet spot” where visitors exercise enough autonomy to express desires but listen to the tour guide enough to benefit from designer intent.

Related Work: Drama Management

We consider tours in a museum to be similar to plot progression in interactive drama. Using a drama manager to guide IE was first proposed in 1986 (Laurel 1986), formalizing the idea of an agent directing the action in response to visitor’s actions. TTD-MDPs were inspired by a formalism for drama management proposed by Bates (Bates 1992), and later formulated as a search problem (Weyhrauch 1997) based on an expecti-max game tree search over plot point sequences, and then reformulated as a reinforcement learning problem (Nelson *et al.* 2006). That work motivated the development of TTD-MDPs for controlling variety of experience.

In each of those three approaches (search, reinforcement learning, and TTD-MDPs) the drama manager acts to optimize an author-supplied evaluation function. Instead, we take a prototype-based approach. Others have used prototypes of desired experiences to guide drama managers (Magerko 2005; Riedl & Stern 2006). One representative approach is the IDA architecture (Magerko & Laird 2004). In IDA, the “director” reasons about future player behavior that might threaten plot progression and acts to explicitly prevent it. This relies on a player that always follows director actions (*i.e.* the only autonomy comes when the guide chooses not to act). In our experiments, we have shown that when the guide’s suggestions are always followed “bad” parts of the space are avoided but player’s tend not to realize their own goals.

Conclusion

We have presented a novel approach to prototype-based drama management that balances player autonomy and authorial intent through the consideration of prototype neighborhoods as well as considers player interaction in a multiplayer environment. In particular, our autonomous tour guides dynamically construct tours online in response to visitors’ reactions to their suggestions in much the same way a drama manager takes actions in a game world to guide a player according to authorial intent.

We are able to show that when visitors cooperate, even only occasionally, they achieve about as many goals as visitors who never cooperate, while significantly reducing congestion and thus increasing everyone’s overall quality of experience. On the other hand, we found that our tour guides perform best when visitors occasionally ignore suggestions. That is, we find that when players exercise their autonomy, ignoring drama manager actions some (but not all) of the time, they are able to realize many of their goals while frequently avoiding potentially bad plot sequences. We have illustrated through a case study that there is a sweet spot between fully author-controlled games and open world games with complete player autonomy.

Acknowledgments

The authors thank Mark J. Nelson for his comments. This research was performed while on appointment as a U.S. Department of Homeland Security (DHS) Fellow under the DHS Scholarship and Fellowship Program, a program administered by the Oak Ridge Institute for Science and Education (ORISE) for DHS through an interagency agreement with the U.S Department of Energy (DOE). ORISE is managed by Oak Ridge Associated Universities under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the authors’ and do not necessarily reflect the policies and views of DHS, DOE, or ORISE. We also acknowledge DARPA support under contract No. HR0011-04-1-0049.

References

Bates, J. 1992. Virtual reality, art, and entertainment. *Presence: The Journal of Teleoperators and Virtual Environments* 2(1):133–138.

Laurel, B. 1986. *Toward the Design of a Computer-Based Interactive Fantasy System*. Ph.D. Dissertation, Drama department, Ohio State University.

Levenshtein, V. I. 1965. Binary codes capable of correcting deletions, insertions and reversals. *Doklady Akademii Nauk SSSR* 163(4):845–848.

Levenshtein, V. I. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady* 10(8):707–710.

Magerko, B., and Laird, J. E. 2004. Mediating the tension between plot and interaction. In *AAAI Workshop Series: Challenges in Game Artificial Intelligence*.

Magerko, B. 2005. Story representation and interactive drama. In *Proceedings of the First Annual Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-05)*.

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 (AAMAS-06)*.

Riedl, M. O., and Stern, A. 2006. 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)*.

Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; Mateas, M.; and Littman, M. L. 2006. Targetting specific distributions of trajecotries in MDPs. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*.

Weyhrauch, P. 1997. *Guiding Interactive Drama*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. Technical Report CMU-CS-97-109.