Creating Customized Virtual Experiences by Leveraging Human Creative Effort: A Desideratum

Mark O. Riedl and Boyang Li
School of Interactive Computing; College of Computing
Georgia Institute of Techology
{riedl, boyangli}@gatech.edu

ABSTRACT

The task of entertaining people has, until very recently, been the exclusive domain of humans. However, recent work in the area of computational creativity, story generation, interactive storytelling, and autonomous believable agents suggests that AI may be used to create dynamic, interactive, and engaging real-time entertainment experiences. In this paper we consider the role of a novel technique called *Experience Adaptation* in the process of creating and delivering customized entertaining experiences. Experience Adaptation is an offline process that leverages human creative ability by taking human-authored storylines – in this case specifications of desired future experiences – and autonomously "re-writing" them based on unique requirements of individual users. With Experience Adaptation, we are working toward effectively scaling up entertainment computing.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems – *games, industrial automation.*

General Terms

Algorithms, Design

Keywords

Experience Adaptation, Narrative Intelligence.

1. INTRODUCTION

Artificial intelligence has long been used to automate certain tasks in order to perform those tasks faster, more accurately, more efficiently, more safely, or more often. However, the task of entertaining people has, until very recently, been the exclusive domain of humans. When it comes to commercial production of entertainment artifacts – TV shows, movies, novels, theatre, computer games, etc. – the task of entertaining people has been the exclusive domain of "creative professions" such as writers, actors, movie directors, theatre and improv performers, dungeon masters, and so on. The reason the task of entertaining people has been the exclusive domain of humans is because the creativity and intuition that human entertainers possess have not been reliably replicated in computational systems.

Currently, there are fewer human "producers" of entertainment than there are human "consumers" of entertainment. This model works fine for mass-consumption entertainment such as film, TV, books, and, to a lesser extent, theatre performances. The *creative authoring bottleneck* refers to the situation where the cost of employing enough professional human producers to satisfy the demands of human consumers is prohibitively high, resulting in a situation where there is more demand for quality content than production of quality content. (We use "authoring" to mean the deliberate creation of any entertainment-related artifact, including an improvised performance created in real-time [5]). Recent work in the area of computational creativity, story generation, interactive storytelling, and autonomous believable agents lays the groundwork for a future where entertainment is fully automated. We are now at a unique point where modern computer technology, simulation, and computer games have opened up the possibility of that more can be done in the area of *on-demand* and *just-in-time* entertainment.

On-demand entertainment refers to the possibility that one can request, at any time, an entertainment experience that is significantly different from any previously consumed. For example, game players can exhaust game-play content faster than expansion packs and new releases can be produced. For an early case study in which consumers outpace producers of content in online virtual game worlds, see [8]. Ideally, there is a one-to-one relationship between producers and consumers so that content can never be consumed faster than it is produced.

Just-in-time entertainment means that entertainment artifacts should be customized or configured based on information that is only available just before it is needed. Just-in-time entertainment affords creation of unique entertainment experience based on a consumer's their needs, wants, ability, and history, all of which cannot be known *a priori*.

As we approach a world in which on-demand and just-in-time entertainment is the expectation, the conventional consumer-producer model breaks down. To overcome the creative authoring bottleneck, we must consider automation. In entertainment, automation is necessary whenever one would want a real person to do *X* in an interactive experience, but sufficient other people are not available in that role. The following are examples of *X*:

- NPCs (shopkeepers, farmers, victims) in computer games.
- Opponents and companions in computer games
- Dungeon master
- Storywriter for books, movies, and games
- Game designer

As we go down this list, an autonomous system is charged with taking progressively more responsibility for the human user's quality of experience. The decisions being made can only be made in a just-in-time fashion because we need to know (a) who the user is, (b) what the user needs, preferences, and desires are, and (c) what the user is doing at any given moment.

While this motivates the need for autonomous systems capable of creativity and expressivity, achieving the goal of autonomous systems capable of assuming responsibility for human users' entertainment experiences is largely an open research question.

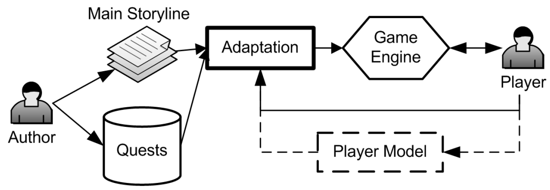

Figure 1. The experience adaptation pipeline.

Until we have computational systems capable of creativity rivaling that of human creators, there is value in exploring *hybrid* approaches in which humans and computational systems share the responsibility of managing human users' entertainment experiences. The goal of such a system is to scale up the human creator's ability to produce meaningful, customized, and potentially highly interactive experiences. We can consider both online and offline approaches. An online system leverages human creativity through semi-autonomy or interpreting and elaborating on human intent (c.f., [3]). An offline system attempts to autonomously modify and/or customize human-authored content to prepare it for real-time interactive experience.

In this paper, we describe an offline system for leveraging human creativity for the purposes of scaling up the production and delivery of unique, customized entertainment experiences. Specifically, we introduce *Experience Adaptation*, as a computational technique that takes a creative authored artifact and adapts it to the specific requirements of a single individual. Experience Adaptation, when implemented for many people – consumers – results in the multiplication of a single creative artifact into many unique creative artifacts.

2. EXPERIENCE ADAPTATION

Considering the creative authoring bottleneck, how do we scale up a human creator's ability to deliver unique, customized, interactive experiences to a large number of consumers of entertainment artifacts? To put it another way: how do we increase *authorial leverage* [1], the ratio of quality of experience to authorial input. Chen et al. [1] measure authorial leverage as the quality of experience per unit of domain engineering, where quality is a function of complexity, ease of change, and variability of experience. Riedl et al. [10] calculate leverage as the number of distinct experiences per unit of domain engineering. In both cases, the number of distinct experiences that can be produced from a static amount of authored content is measured.

We propose a technique whereby a few human-authored descriptions of experiences to be had in a virtual world are leveraged to provide numerous experiences customized to individuals. The technique, *Experience Adaptation*, computationally takes a single, human authored story and autonomously customizes it to individuals' unique needs, wants, and desires. Experience Adaptation leverages the creative abilities of a single human author into multiple playable experiences.

The Experience Adaptation pipeline is shown in Figure 1. A human author develops a storyline as a means of describing what a user should experience in the virtual world. The storyline determines events that will happen in the virtual world, including specifications for the behaviors of non-player characters. The storyline, provided in a computational format that facilitates automated analysis and reasoning, is combined with a player model and a world model. The world model describes what characters – human or virtual characters – can do in the world, and how the world is changed when actions are performed. The player model provides information about the user in terms of preferences over experiences. The player model also contains historical information describing the types of experiences the user has previously had. The player model is capable of generating a set of experiential requirements – the features of the experience the user should receive. See Medler [6] and Thue et al. [13] for perspectives on player modeling. Currently, we allow the player to directly specify what he or she desires in an experience.

The storyline, player model requirements, and world model are inputs into the Experience Adaptation system. The storyline is analyzed to determine whether it meets the experiential requirements from the player model. If it does not, the Experience Adaptor engages in an iterative process of making changes to the storyline until it meets the requirements of the user model. The result is a new creative artifact describing a customized story experience, which is sent to an appropriate game engine for interactive real-time execution. Note the cycle in Figure 1 created by the Experience Adaptation process, resulting in greatly improved replayability of authored experiences; as the player model evolves over time, the same human-authored storyline can be recycled into unique experiences.

The core component in the Experience Adaptation process is the Experience Adaptor. The Experience Adaptor has two functions, to interpret the requirements provided by the user model, and to "rewrite" the story provided by a human author. The *Experience Adaptation Problem* is as follows: given a domain model, a set of experiential requirements, and a storyline that does not meet the requirements, find a coherent storyline that meets the experiential requirements and preserves the maximal amount of original content. A *coherent* storyline is one in which all events have causal relevance to the outcomes [14]. The preservation of original content ensures that as much of the creative intuition of the human author remains intact as possible.

The storyline can be adapted in three different ways:

- **Deletion** – Events in the storyline can be removed because they are unnecessary or unwanted.
- **Addition** – Events can be added to the storyline to achieve experiential requirements, and to ensure narrative coherence.
- **Replacement** – a combination of deletion and addition, old events are swapped for new events that better achieve experiential requirements.

The application of these operations enables a refinement-search algorithm to incrementally tear down and build up a complete, human-authored narrative structure until it meets the experiential requirements.

2.1 Computational Representation of Experience

We represent experience as a narrative. A narrative is a sequence of events with continuant subject and constitutes a whole. In this case, the narrative is a description of the *expected* sequence of events that will occur in a virtual environment. Computationally, we represent narratives as a partial-order plan, which provides a formal framework for which to reason about changes to narrative built on first principles (for example, we can ask if a narrative is sound). Specifically, we employ a specialized plan representation

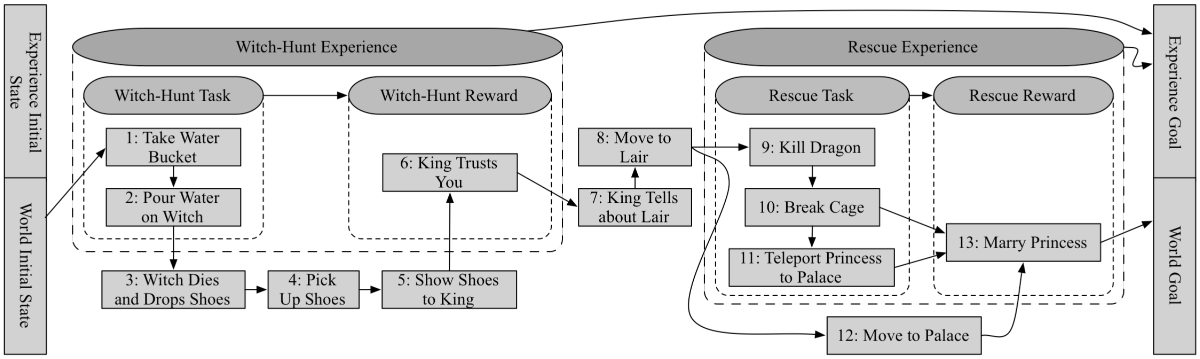

Figure 2. The original experience storyline.

from Decompositional Partial Order Planning (DPOP) [15], a combination of partial-order planning and hierarchical task network planning. In a DPOP plan, actions are related via causal links, temporal constraints, and decompositions. A causal link, denoted $a_1 \rightarrow^c a_2$, specifies that action a_1 established a condition c in the world that is causally necessary for latter action a_2 to occur. Temporal constraints determine if one action must strictly occur before another. Decompositions relate abstract actions to sets of less abstract actions.

By representing narratives as plans, actions indicate events that are expected to happen during the user's interactive experience. Figure 2 shows an example of a DPOP plan representing an experience one might have in a role-playing game. Primitive world-level events are shown as boxes and abstract events are shown as arrows. Arrows represent causal links. Not all causal relations are shown. Dashed lines encapsulate events that make up decompositions. The events that are essential for each experience occur within the hierarchical decompositions, while other events are considered incidental. The events are numbered in chronological order.

The initial state is the description of what the author assumes the state of the virtual world would be like before the experience begins and the goal situation describes the way the author expects the user and the virtual world will be different after the experience terminates. The goal situation is partially in terms derived from the experiential requirements. Thus, one of the responsibilities of the Experience Adaptor is to reconcile the initial state and goal situation with just-in-time information about the user.

2.2 The Adaptation Process

The adaptation process involves two steps: (a) storyline analysis, and (b) storyline reconciliation. Storyline analysis is a process whereby the experiential requirements from the player model are compared to the storyline. In the case that there is some aspect of the storyline that does not meet the set of experiential requirements, the storyline analysis begins deleting undesired events and updating the plan goal and initial state. Typically, this involves deleting an abstract experiential event and all of its children. In the process of deleting experiential and world events, storyline analysis causes numerous inconsistencies in the plan. Inconsistencies include gaps in the plan due to deleted events, unsatisfied goals, and mismatches between the events in the plan and the initial state.

Storyline reconciliation is a refinement search process in which inconsistencies are eliminated. The refinement search algorithm is an extension of DPOP. The Adaptation algorithm is specifically designed to address the following challenges:

1. The starting point of the planning process is a partially complete plan instead of the typical empty plan. This is significant because existing events in the initial plan can cause aesthetically unappealing results or outright failure. Thus the Adaptation algorithm has the option of removing events that were hand-authored.
2. Causal coherence is maintained. *Coherence* means that all events have causal relevance to the outcome. Planners guarantee that all events' preconditions are satisfied; there is a causal path from initial state to every action. Cohesion and coherence require that all events have effects that lead, in a significant way, to the goal state. An event that does not have an effect that leads to the goal state is a *dead end*, which has a negative impact on one's comprehension of the narrative structure [14] and, we believe, one's satisfaction with the experience.

The Adaptation algorithm, thus, works as follows. Starting with the original storyline plan with unwanted events deleted and initial state and goals updated, flaws are identified. Flaws include the standard set of DPOP flaws: an abstract action that is not decomposed; an action has a precondition that is not satisfied by a causal link, and causal threats (c.f. Young and Pollack [15]). These flaws are resolved in the standard ways: by instantiating new events, adding causal links between existing and new events, or asserting temporal orderings. Due to the fact that we are starting with a complete plan, we also allow the algorithm to resolve unlinked precondition flaws and causal threats by removing the action in question.

In addition, we introduce a *dead-end flaw* indicating that an event does not have an effect that contributes to a causal chain leading to a goal. Dead-end flaws occur because we are working from an existing plan structure with gaps; during the storyline analysis stage, causal links can be broken. Conventional DPOP only considers whether it is possible for an event can occur, by establishing a causal chain from the initial state to the event. From a storytelling and experiential perspective, we also require all events to contribute to the achievement of some meaningful goal situation. Dead-end flaws can be repaired with the following strategies:

- Extend a causal link from one of the effects of the dead-end event to an unsatisfied precondition of another event.
- Move the initiation point of an existing causal link from

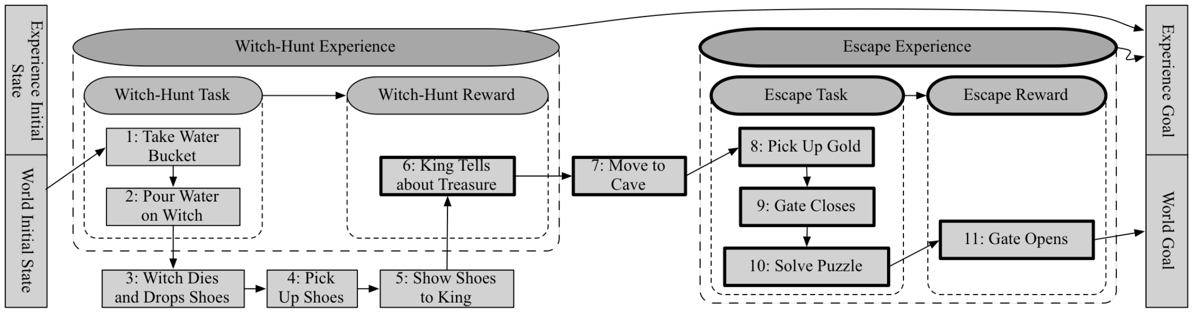

Figure 3. The adapted experience storyline, replacing the princess rescue and marry events with a treasure hunt.

some event to the dead-end event. Note that this may make the other event a dead-end.

- Remove the dead end event.
- Do nothing, leaving the final result non-coherent.

Being able to further remove events from the plan structure ensures that no event from the original plan can cause the plan refinement process to fail. This can happen if the effects of the event prevent other actions from being inserted. Note that to preserve systematicity, the algorithm can only remove events that came from the original narrative plan, preventing infinite loops of addition and removal. Most significantly, removal of events in dead-ends means that original events that cannot be used in the adapted plan are not kept around because they could look odd, unnecessary, or irrelevant.

2.3 Example

To provide a motivating example, consider the short storyline in Figure 2 meant to be played as an interactive role-playing game. The background is that the player is an adventurer who has heard of a feud between the king and a witch. The storyline shows a sequence of events in which the player kills the witch, gains the trust of the king, and is sent to rescue the princess. The sequence culminates in the player marrying the princess. However, suppose the player is not interested in rescuing and marrying a princess. Instead, the player is motivated by material achievement, such as acquiring gold. Storyline analysis results in the removal of the rescue experience. The goal is updated to require that the player experience an escape and, on the world-level, that the player acquires wealth.

The Experience Adaptor revises the storyline to make the witch hunt and escape work together seamlessly. The adapted storyline is shown in Figure 3, with changes in bold. Specifically, the escape experience is inserted into the plan. The events in which the player learns about and moves to the lair become dead-ends and are removed. The event in which the king trusts you also becomes a dead-end and is replaced by a new reward, being told about the treasure, which is linked into previously existing causal chain. Finally, to link the witch-hunt experience to the escape experience, an event in which the player moves to the cave is inserted into the storyline.

2.4 Authoring and Scaling

The authoring process is as follows. First, there must be a world domain model, containing specifications for primitive event actions. Second, some number of experience fragments must be authored as DPOP recipes. These first two steps constitute a one-time authoring cost by a domain engineer. Next, one or more storylines may be authored in the DPOP representation such that they consist of experience fragments and other primitive events that connect the fragments. This may require additional effort on the part of the human author, but the payoff for this extra effort is an exponential scaling of the initial effort.

Theoretically, experience adaptation takes a single storyline and produces as many adaptations as the size of the power set of experience fragments. In practice, the number of pragmatic adaptations will be lower because it's likely that a large fraction of the original is retained in each adaptation request. However, the scaling will still be exponential.

To *manually* achieve this scaling, one would have to author $n(n-1)$ experience fragments (n-1 variations of each experience so it can be paired with n-1 other experiences) and use a simple algorithm for appropriately selecting and sequencing those $n(n-1)$ fragments. Under the manual scheme, adding new experience fragments becomes increasingly difficult. With Experience Adaptation, new experiences can be added independently of any other in the library, assuming a sufficiently rich world domain model. One benefit of Experience Adaptation that cannot be reproduced manually comes in the ability to arbitrarily make changes to the initial state and goal situation to create fine-grain adaptations such as swapping a magic potion for gold, or adjusting difficulty by making the witch immune to water.

Our current prototype Experience Adaptor has been tested in the context of role-playing game storylines (such as that in Figure 2) and on military training scenarios. Future work is required to measure the pragmatic authorial leverage [1, 10] of the system in terms of authoring effort versus effective output. An evaluation of aesthetic quality of generated storylines is currently underway.

3. RELATED WORK

Experience Adaptation is being simultaneously explored in the context of military training, under the moniker *Scenario Adaptation* [9]. In Scenario Adaptation, military training scenarios are represented as DPOP plans, which are adapted in order to keep the learner in his or her zone of proximal development.

Thue et al. [13] describe a technique whereby a player model based on role player types is used to select branches through an interactive story. This approach assumes the existence of a branching story graph. Hullett and Mateas [2] describe a technique whereby experiences of users are changed by reconfiguring the level. Experience is equated with navigation through a virtual space. Experience Adaptation will also require coordination between storyline and the virtual environment and future work will likely involve some degree of automated reconfiguration of a virtual world.

Experience Adaptation is *not* Interactive Storytelling. Interactive Storytelling systems demonstrate how players or learners may interact with story and scenario content in complex simulation environments. Typically, an intelligent agent called a *Drama Manager* adjusts the virtual environment – including the behaviors of virtual characters – during execution to meet dramatic or learning objectives. See Roberts and Isbell [12] and Riedl et al. [10] for an overview of interactive storytelling and drama management systems. The distinction between Experience Adaptation and Interactive Storytelling is that in Interactive Storytelling adjustments to the virtual world occur at execution time in order to cope with the real-time actions of the player. Experience Adaptation, on the other hand, "rewrites" the objectives of the virtual environment in an offline process. In this light, Experience Adaptation and Drama Management are complimentary: the Experience Adaptor configures the Drama Manager, which oversees the user's interactive experience.

The Drama Management technique known as *Narrative Mediation* is especially relevant. The Automated Story Director framework [10] in particular makes partially ordered plans interactive by generating branches and rendering storyline events into goals that dictate the behaviors of semi-autonomous character agents. We envision systems such as this can be used in conjunction with Experience Adaptation to deliver highly individualized, interactive experiences.

As an offline procedure, Experience Adaptation is a form of story generation. Story generation is the process of automatically creating novel narrative sequences from a set of specifications. The most relevant story generation work is that that uses planning as the underlying mechanism for selecting and instantiating narrative events (c.f., [7, 4, 11]). The distinction between our Experience Adaptor and story planning is that the Experience Adaptor starts with a complete, sound narrative structure and is capable of *removing* events.

4. CONCLUSIONS

This paper addresses the problem of leveraging human-authored content in order to scale-up the applicability of virtual entertainment experiences. We assume that experiences are narratives that describe how the experience is expected to unfold. We introduce the concept of employing an offline process to customize experience specifications. Specifically, a player model determines the desired features of a real-time experience and an Experience Adaptor automatically customizes a hand-authored storyline accordingly. To that end, we draw heavily from recent work on search-based narrative generation, although adaptation is necessarily different due to the fact that it starts with a complete, human-authored narrative.

While this paper focuses on offline aspects of Experience Adaptation, the process is in service of creating compelling interactive, real-time experiences. Interactive real-time experiences require autonomous believable characters and the ability to dynamically adapt the storyline to accommodate the user's moment-to-moment decisions. Future work considers the use of interactive narrative techniques such as that of Riedl et al. [10], which uses a combination of dynamic narrative re-planning and semi-autonomous character agents to create a real-time experience. Ultimately, we believe that a combination of offline and online AI processes will be required to solve the general problem of customizing interactive entertainment experiences. As on-demand and just-in-time entertainment computing becomes reality, the greater the need for autonomous systems capable of creativity and expressivity.

5. REFERENCES

[1] Chen, S., Nelson, M.J., Sullivan, A., and Mateas, M. 2009. Evaluating the Authorial Leverage of Drama Management. In Proceedings of the AAAI Spring Symposium on Intelligent Narrative Technologies II.

[2] Hullett, K. and Mateas, M. 2009. Scenario Generation for Emergency Rescue Training Games. In Proceedings of the Fourth International Conference on the Foundations of Digital Games.

[3] Ishigaki, S. White, T., Zordan, V., and Liu, K. 2009. Performance-Based Control Interface for Character Animation. In Proceedings of SIGGRAPH'09.

[4] Lebowitz, M. 1987. Planning stories. In Proceedings of the 9th Annual Conference of the Cognitive Science Society.

[5] Magerko, B.S. and Riedl, M.O. 2008. What Happens Next?: Toward an Empirical Investigation of Improvisational Theatre. In Proceedings of the 5th International Joint Workshop on Computational Creativity.

[6] Medler, B. 2009. Using Recommendation Systems to Adapt Gameplay. International Journal of Games and Computer-Mediated Simulations 1(3).

[7] Meehan, J. 1976. The Metanovel: Writing Stories by Computer. Ph.D. Dissertation, Yale University.

[8] Morningstar, C. and Farmer, F.R. 1991. The Lessons of Lucasfilm's Habitat. In Cyberspace: First Steps, M. Benedikt, Ed. MIT Press.

[9] Niehaus, J. and Riedl, M.O. 2009. Scenario Adaptation: An Approach to Customizing Computer-Based Training Games and Simulations. In Proceedings of the AIED 2009 Workshop on Intelligent Educational Games.

[10] Riedl, M.O., Stern, A., Dini, D., and Alderman, J. 2008. Dynamic Experience Management in Virtual Worlds for Entertainment, Education, and Training. International Transactions on Systems Science and Applications 3(1).

[11] Riedl, M.O and Young, R.M. 2004. An Intent-Driven Planner for Multi-Agent Story Generation. Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems.

[12] 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 3(1).

[13] Thue, D., Bulitko, V., Spetch, M., and Wasylishen, E. 2007. In Proceedings of the 3rd Artificial Intelligence and Interactive Digital Entertainment conference.

[14] Trabasso, T. and van den Broek, P. 1985. Causal Thinking and the Representation of Narrative Events. Journal of Memory and Language 24.

[15] Young, R.M. and Pollack, M.E. 1994. Decomposition and causality in partial-order planning. In Proceedings of the 2nd International Conference on Artificial Intelligence and Planning Systems.