

# Interactive Narrative: An Intelligent Systems Approach

Mark O. Riedl<sup>1</sup> and Vadim Bulitko<sup>2</sup>

<sup>1</sup>School of Interactive Computing, Georgia Institute of Technology

<sup>2</sup>Department of Computing Science, University of Alberta

## Abstract

Interactive narrative is a form of digital interactive experience in which users create or influence a dramatic storyline through their actions. The goal of an interactive narrative system is to immerse the user in a virtual world such that he or she believes that they are an integral part of an unfolding story and that their actions can significantly alter the direction and/or outcome of the story. In this article we review the ways in which artificial intelligence can be brought to bear on the creation of interactive narrative systems. We lay out the landscape of about 20 years of interactive narrative research and explore the successes as well as open research questions pertaining to the novel use of computational narrative intelligence in the pursuit of entertainment, education, and training.

## 1 Introduction

Storytelling, in oral, visual, or written forms, plays a central role in various types of media, including novels, movies, television, and theatre. The prevalence of storytelling in human culture may be explained by the use of narrative as a cognitive tool for situated understanding [3]. This *narrative intelligence*—ability to organize experience into narrative form—is central to the cognitive processes employed across a range of experiences, from entertainment to active learning. It follows that computational systems possessing narrative intelligence may be able to interact with human users naturally because they understand collaborative contexts as emerging narrative and are able to express themselves by telling stories. A number of narrative intelligence tasks have been studied from a computational perspective including story understanding, story generation, and commonsense reasoning.

One of the most compelling applications of narrative intelligence is the prospect of *interactive narrative*. Interactive narrative is a form of digital interactive experience in which users create or influence a dramatic storyline through actions, either by assuming the role of a character in a fictional virtual world, issuing commands to computer-controlled characters, or directly manipulating the fictional world state. It is most often considered as a form of interactive entertainment, but can also be used for serious applications such as education and training. The most common form of interactive narrative involves the user taking on the role of the protagonist in an unfolding storyline. The user can also be a disembodied observer—as if watching a movie—but capable of making changes to the world or talking to the characters. The goal of interactive narrative is thus to immerse the user in a virtual world such that he or she believes that they are an integral part of an unfolding story and that their actions have *meaningful* consequences. That is, the user’s actions can be observed to have a direct impact on the direction or outcome of the storyline. There are many digital entertainment applications, such as computer games, that use stories to structure users’ activities. The distinction between interactive narrative and other forms of digital entertainment is that interactive narrative systems afford the player to act in ways that fundamentally alter the direction and/or outcome of the unfolding storyline.



Figure 1: The *Holodeck* from an episode of *Star Trek: The Next Generation*.

To illustrate the significance of a user able to act to change the direction and/or outcome of a narrative as it is unfolding, consider the *Holodeck* from the popular television series, *Star Trek: The Next Generation* (see Figure 1). The *Holodeck* is a fictitious technology that uses holography to immerse humans in a photorealistic virtual reality, often populated by virtual, intelligent, computer-controlled characters. While the *Holodeck* is portrayed as having many serious uses, one of its primary uses is entertainment via the immersion of its users in fictional worlds. Examples from the TV series range from Wild West to Film-Noir to Victorian dramas. Suppose one were to enter the Holodeck to step into the shoes of Shakespeare’s *Hamlet*.<sup>1</sup> Further suppose that the user chooses not to reenact the actions and dialogue of the Hamlet character, who remains locked in indecision throughout a majority of the story. Instead, suppose the user strikes out in a new direction, perhaps confronting and slaying the antagonist, Claudius, in the first act instead of the third. What should the other characters do in response? What should happen next? Can the story even continue? Shakespeare has nothing to tell us about how to proceed.

Setting aside the questions of graphics rendering, holography, and virtual reality—stories are not only told in visual media—the example above presents some intriguing research questions:

1. How does a computational system determine what should happen next in the narrative?
2. What does a computational system need to know about the fictional world and the user in order to reason about the user’s narrative experience?
3. How does a computational system intervene in the fictional world to bring about an engaging experience?

The first two questions pertain to the construction of computational systems that exhibit narrative intelligence. The answers to these questions not only take us one step closer to automated creation of engaging experiences in virtual worlds, but also addresses the fundamental quest for intelligent systems that exhibit human-level capabilities; the ability to create stories is an ability that is, to date, unique to humans. The third question pertains to the construction of computational systems that can act on behalf of the user to his or her benefit, be it entertainment, education, or training.

Interactive narrative has its roots in entertainment. However, a small but growing faction of researchers are considering the role of interactive narrative in “serious” application such as education, training, advertising, and argumentation. In science, technology, engineering and mathematics (STEM) education, interactive narrative shares many parallels with problem-based inquiry learning, which involves guiding a

---

<sup>1</sup>The example of the recreation of *Hamlet* in a *Holodeck* was inspired by Murray [8]. Cavazza and colleagues [10] have directly experimented with an interactive version of Shakespeare’s *The Merchant of Venice*.

student through a succession of problems that build off one another (cf., [13]). Procedural skill learning often employs interactive scenarios to create realistic contexts for skill practice [11].

## 2 Experience Management

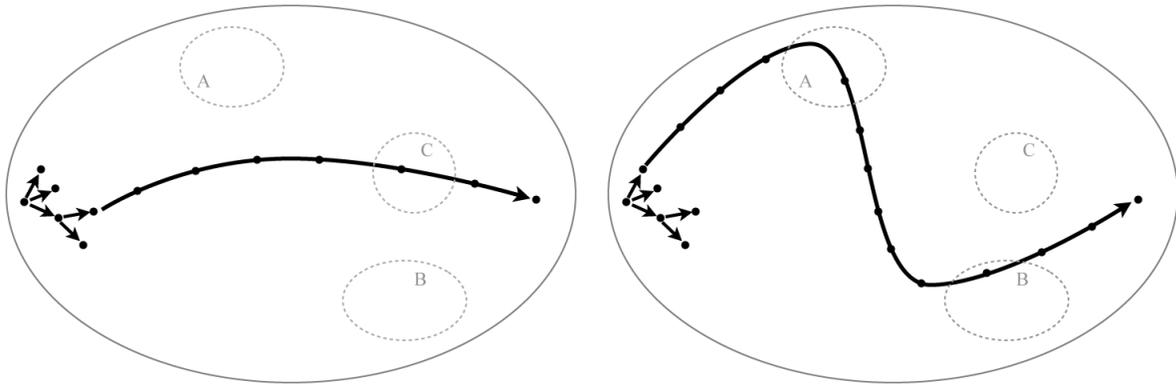
The core research challenge is how to balance the need for a coherent story progression with user agency, which are often at odds. A coherent narrative experience is one in which all events build off prior events until a conclusion is reached. The user may act in a manner he or she deems best, unintentionally introducing inconsistencies—events that the system cannot build off—or making it impossible for subsequent events to unfold as planned. Typically this occurs because the user is unaware of the ways in which the narrative may unfold, but may also occur because the user is trying to steer the narrative in a new direction or actively and mischievously testing the bounds of the system’s responsiveness. The key challenge to interactive narrative is how to balance these competing needs to ensure the player feels he or she has agency to affect the direction and/or outcome of his or her narrative experience while still ensuring the experience is coherent.

A common solution, first proposed by Bates [2] is to implement a *drama manager*. A drama manager is an intelligent, omniscient, and disembodied agent that monitors the virtual world and intervenes to drive the narrative forward according to some model of quality of experience. An *experience manager* is a generalization of this concept, recognizing the fact that not all narratives need to be dramatic, such as in the case of education or training applications. An experience manager drives the narrative forward by intervening in the fictional world, typically by directing computer-controlled characters (called *non-player characters* (NPCs)) in how to respond to the user’s actions. To that end the user should not be aware of the existence of the experience manager or its interventions.

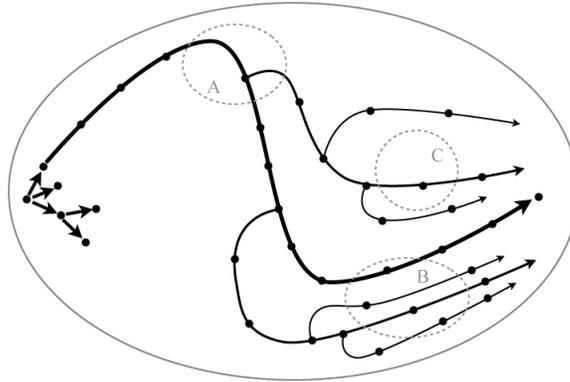
How does an experience manager know how and when to intervene in the virtual world? An experience manager must generally look ahead into possible futures of the user’s experience to determine the best intervention, if any, to bring about a structurally coherent experience. Unlike activity recognition or prediction, the experience manager seeks the best *narrative* sequence according to the narratological principles of coherence and other criteria for experiential quality. Because of the experience managers ability to intervene through NPCs, this sequence does not have to be—and often should *not* be—the most likely sequence of events. Instead, the future narrative should be the best narrative according to some criteria that can be achieved given the agency of the user and the ability to intervene. The projection of a narrative sequence into the future enables the experience manager to evaluate the global structure of possible player experiences in a way that cannot be achieved by looking at any single world state in isolation. Armed with this knowledge, the experience manager must reason about the effects of its interventions in the virtual world to bring about the desired narrative experience.

Looking at experience management from the perspective of states and actions, there are a number of states the fictional world can be in, and a number of actions that can be performed by the user, by NPCs, and directly by the experience manager itself. Consider the simple example illustrated in Figure 2. The large oval represents the set of all states that the virtual world can be in. Starting with an initial state on the left, the user and NPCs can perform actions that transition the virtual world into new states. Suppose the user and computer-controlled characters were left to their own devices and the resultant state space trajectory is shown in Figure 2(a). Further suppose that from a narratological perspective, any trajectory that passes through a state in C will be considered a “poor” experience by some criteria. Meanwhile, a narrative that passes through a state in A followed by a state in B will be considered a “good” experience by the same criteria. The goal of the experience manager is to intervene in such a way that the user’s experience follows a trajectory more similar to that shown in Figure 2(b). Of course, as the user has opportunities to act, the experience manager must continuously project and select possible future narrative trajectories, as in Figure 2(c).

How does an experience manager determine whether a given intervention will increase the quality of the user’s experience? In theory, an experience manager is capable of reasoning about the quality of



(a) A possible narrative trajectory through state space. (b) A possible narrative trajectory that visits states deemed favorable and avoid states deemed unfavorable.



(c) Accounting for player interaction.

Figure 2: The experience management problem is to compute trajectories through state space.

the user’s experience, including but not limited to: hard or soft constraints on narrative structure (as above), principles of dramatic tension, pedagogical goals, or other aesthetic considerations. These criteria, principles, goals, and aesthetic considerations are provided by a *human author*. The human author is the person behind the scenes that desires to tell a story in an interactive fashion but cannot be present at execution time to shape the user’s experience his or herself.<sup>2</sup> The relationship between the human author and experience manager is of critical concern to interactive narrative research, as the experience manager can be considered an autonomous surrogate for the human author.

### 3 A Taxonomy of Approaches to Interactive Narrative

The history of research and development of interactive narrative systems has been rich and varied. We have mapped the relevant landscape of interactive narrative systems in Figure 3. We organize the landscape into three dimensions that can be used to distinguish different approaches to the problem of creating interactive narrative experiences:

1. **Authorial intent.** To what extent is the human author’s storytelling intent constrains the interactive narrative system? On the left are systems that are highly constrained to carrying out the

<sup>2</sup>Tabletop Role Playing Games such as *Dungeons and Dragons* use a human *game master* or *dungeon master* to drive a plotline forward and to manipulate non-player characters. To that end, experience management is an attempt to automate the role of game/dungeon master either by making decisions that shape the narrative or by carrying out a human game master’s pre-determined decisions at run-time on their behalf.

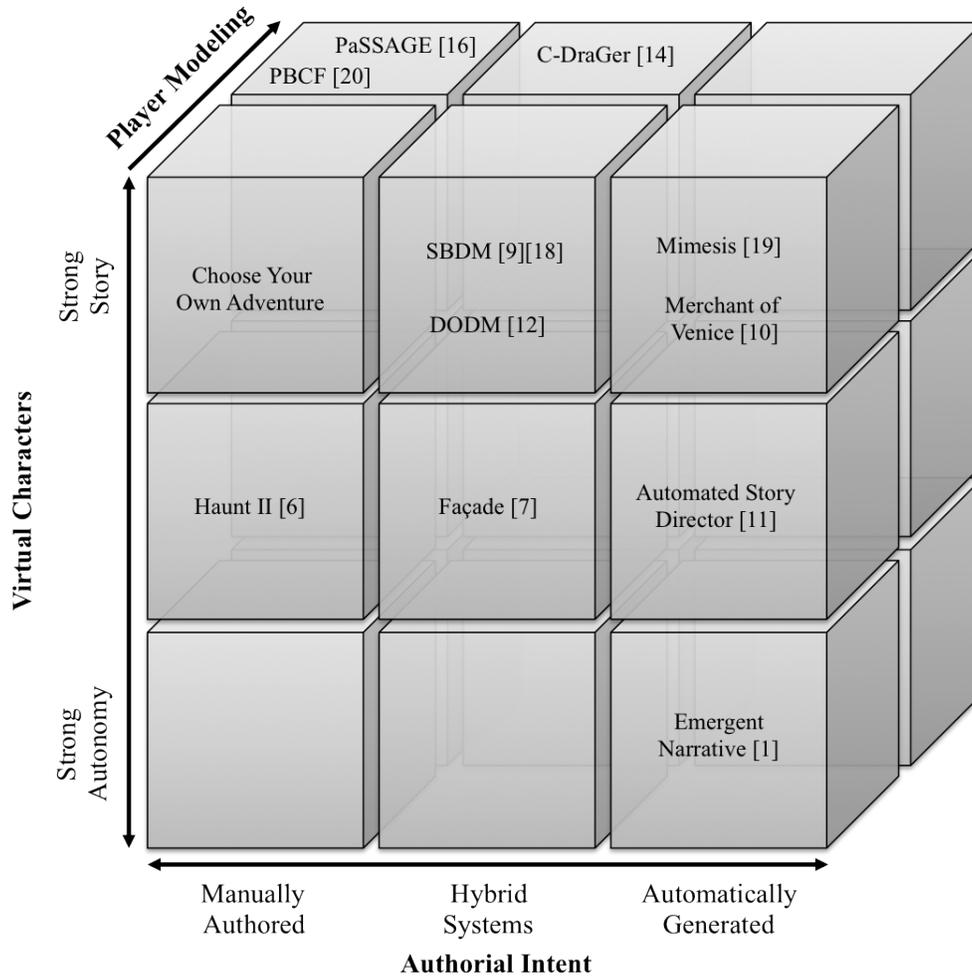


Figure 3: The landscape of interactive narrative research.

human author’s intent. On the right are systems that assume creative responsibility for the user’s narrative experience.

2. **Virtual character autonomy.** The non-player characters in the virtual world can have more or less autonomy from the experience manager. Strong story systems are those in which the NPCs are completely controlled by the experience manager. Strong autonomy systems are those in which the NPCs are unaware of the overarching narrative needs. There is a tension between the needs for NPCs to act consistently with the narrative and the need to act consistently with their own character and settings.
3. **Player modeling.** To what extent does the interactive narrative system attempt to learn about the individual differences of the user? All interactive narratives adapt themselves to the user by observing and responding to the user’s actions. Player models are abstractions over the user’s patterns of behavior in the space of narratives that capture and predict aspects of user behavior that can subsequently be used to manage the experience.

There are too many systems to cite here; we discretize the space and show how a small sampling of systems are situated within this space. In the next sections, we discuss the different dimensions and their implications for the development of interactive narratives.



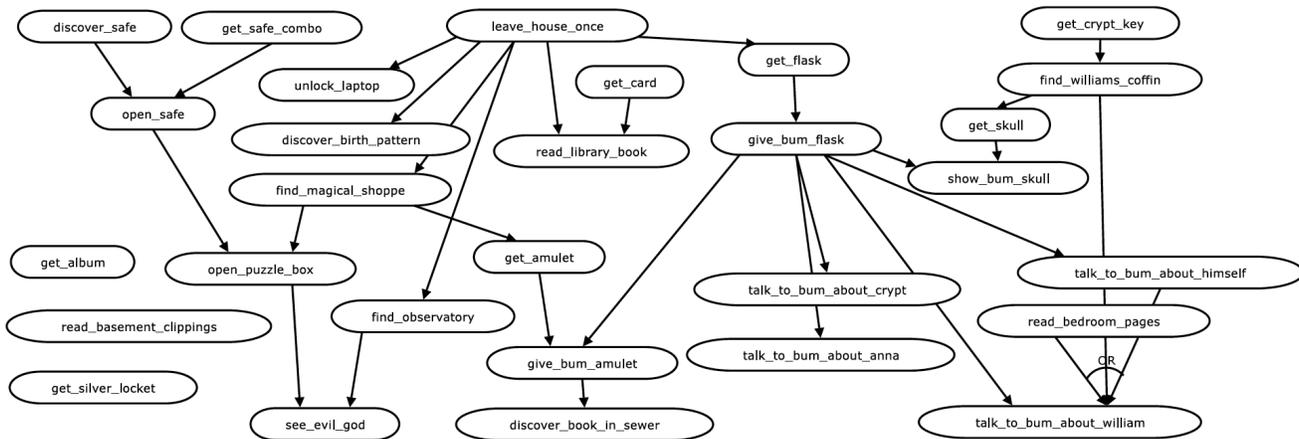


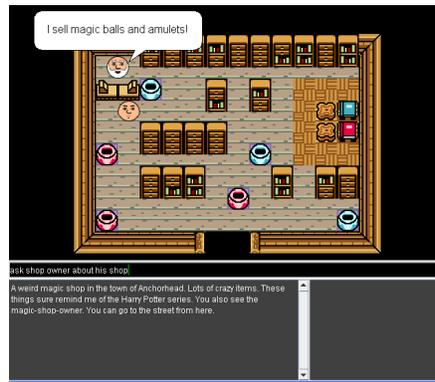
Figure 5: A portion of the plot graph for the interactive fiction, *Anchorhead* [9] (reproduced with permission).

literal specification of branching story graphs. A common approach is to encode this knowledge as a *plot graph* in which nodes are narrative events and arcs denote precedence constraints such that no event can occur unless all events constrained to occur prior to it have already occurred (see Figure 5). Some events are *experience manager actions* (not shown) that can be performed at any time to increase the probability that certain trajectories occur. Examples of experience manager actions include hinting, causing certain plot points to occur, and temporarily denying certain plot points from occurring. A plot graph thus defines a space of possible narrative trajectories; constraints prune out sequences that do not logically make sense or should not be allowed to occur in the user’s narrative experience. A search process—adversarial search [18], reinforcement learning [12], or case-based planning [14]—generates possible trajectories, including experience manager actions, and evaluates the trajectories according to an author-defined heuristic.

As we move farther along the spectrum, the knowledge possessed by the experience manager necessarily becomes more general and the experience manager assumes full responsibility for constructing “good” narratives. The *generative experience management* approach embeds a complete story generator (see [4] for an overview of story generation systems as creative systems) within the experience manager [19, 11, 10]. The knowledge provided by the human author is of much more general form, and the experience manager must generate the narrative from scratch via the embedded story generator. For example, the story generator employed within the experience manager may encode knowledge about the fictional world in the form of a *domain theory*—a declaration of what actions are possible in the fictional world and the world state conditions under which they can be executed without further knowledge about the order that actions should occur.

A generative experience manager must solve the *boundary problem* [6]—to recognize and respond to (intentional or unintentional) attempts by the user to perform actions that deviate from the narrative the experience manager desires to tell. As a simple example, consider the possibility that the user decides to kill an NPC that plays a significant role later in the narrative; the narrative will not be able to continue as expected. Once a deviation is detected, the experience manager can respond in one of two ways:

- **Intervention:** the experience manager can act to prevent the user from crossing the narrative boundary either by directing NPCs to interact with the user in different ways [6] or by changing some aspect of the fictional world (e.g., by causing the user’s gun to jam) [19].
- **Accommodation:** the experience manager allows the deviation to occur and then invokes the story generator to create a new narrative that eliminates any causal inconsistencies that arise from the user’s actions (e.g., by selecting a new NPC to take the role of the killed character) [19, 11].



(a) *Haunt II* [6] attempts to preserve a pre-authored narrative (reproduced with permission). (b) *C-DRaGer* [14] uses a mix of authored knowledge and search (reproduced with permission).



(c) The *Automated Story Director* [11] (left) and an interactive *The Merchant of Venice* [10] (right; reproduced with permission) use generative experience management.

Figure 6: Interactive narrative systems representing different degrees of authorial intent.

Story generation is an open research question and those who pursue generative experience management are interested in the extent to which the story generator or experience manager can assume the creative responsibilities of the human author.

### 3.2 Virtual Character Autonomy

The dimension of virtual character autonomy is concerned with the degree to which computer-controlled entities can act independently of the experience manager, which has implications for *character believability*. A *believable character* is a virtual agent that fosters suspension of disbelief that the user is interacting with a personality-rich, intelligent being (a person, an anthropomorphized animal, etc.) [2]. Believable characters exhibit personality and emotion as they interact in real-time with the environment and the user. Stories are about characters and their interactions; character are thus an integral aspect of interactive narrative systems. However, there can be a tension between an experience manager, which is trying to bring about a global structure to the user's interactive experience, and virtual characters, which are attempting to perform local actions that promote believability. There are a range of approaches to integrating character and story [7]:

- **Strong Story:** Virtual characters do not act without the guidance and permission of the experience manager.



Figure 7: The *Façade* interactive drama [7] strikes a balance between story story and strong autonomy (reproduced with permission).

- **Strong Autonomy:** Each virtual character is a fully autonomous agent, unaware of the needs of the overarching narrative.

Most interactive narrative systems implement an approach to NPCs somewhere between these extremes.

Under the strong story perspective, the experience manager can achieve the highest degree of leverage over the virtual world to bring about the desired narrative experience for the user. Enabling the experience manager to reason about every moment of a virtual character’s interaction or dialogue with the user is generally intractable. Not every detail of every human-NPC interaction or dialogue act is significant to plot progression and an experience manager will reason at the level of abstract units of plot, using processes to decompose plot units to actions when necessary (called *realization*).

Under the strong autonomy perspective, the user’s experience is entirely driven by the uncoordinated decisions of the NPCs and his or her own actions. When NPCs have complete autonomy, an interactive narrative is referred to as an *emergent narrative* [1]. As in the case of simulation games or exploratory learning environments, it is not always necessary to intervene to guide the user’s narrative experience toward a particular conclusion.

Many interactive narrative systems attempt to strike a middle ground between strong story and strong autonomy. The *Façade* interactive drama [7] specifically addresses this issue by using personality-driven decompositions from high-level plot units to executable behaviors. An experience manager provides coherence. Characters are autonomous to the extent that they can independently determine *how* the plot units are realized. The *Automated Story Director* [11] implements *semi-autonomous characters* that act fully autonomously until directed by an experience manager, at which point characters must reason how to seamlessly transition between autonomously selected behaviors and behaviors required for plot progression.

### 3.3 Player Modeling

Up to now, we have made the assumption that the experience manager is a surrogate for the human author by representing his or her authorial intentions in the face of player autonomy. However, different users may have different preferences over how the narrative unfolds. An experience manager may *also* take user’s preferences into account by modeling and predicting the user’s preferences for different narrative futures while attempting to preserve the author’s intent. While all interactive narrative systems adapt the narrative structure to respond to user actions, we use the term *player modeling* to refer to the process of learning a model of the user’s individual differences (e.g., preferences, play style, etc.) While player



Figure 8: The *PaSSAGE* system guides the user down narrative branches that match the user’s preferred style of play.

modeling is becoming more common in computer games, it has only recently been explored in the context of interactive narratives.

One type of player modeling is done by mapping observed player behavior in the fiction world into abstractions. Such abstractions can be later exploited by an experience manager to intervene in the world to the benefit of the user’s perceived experience [16, 14]. Another type of player modeling lies with eliciting structured feedback from users across many interactive narrative experiences [20]. Work on player modeling in interactive narratives is still in early stages and focuses mainly on approaches that strongly conform to authorial intent (see Figure 3); other player modeling techniques exist and are yet to be explored in the context of guiding interactive narrative experiences. Non-player characters can also benefit from player modeling by adapting their behaviors or dialogue according to what they believe about the user.

Personalization of narrative experiences is especially relevant to serious applications of interactive narrative to education and training. In education, an *intelligent tutoring system* (ITS) attempts to emulate the one-on-one interactions between a dedicated human tutor and a pupil. An ITS can be characterized by two processes: (a) determining the best next problem for the learner to work on, and (b) scaffolding of the learner’s problem-solving [17]. ITSs utilize a learner model that captures individual skill and knowledge differences. For inquiry learning and training domains where the problem execution is embedded within a narrative or scenario, a learner-model-driven interactive narrative can scaffold the learner’s problem-solving. If the interactive narrative system employs a *generative experience manager*, it can also be seen as performing problem generation.

## 4 Discussion

There are numerous approaches to interactive narrative systems, not all of which require artificial intelligence. The use of artificial intelligence in the implementation of interactive narrative systems increases the expressive power of the system by (partly) assuming creative responsibility for the user’s narrative experience. This in turn can yield greater responsiveness and range of narratives without diminishing player agency.

Interactive narrative has been an applied area of artificial intelligence and machine learning research

for approximately 20 years. In this time there have been significant advances in the ability for intelligent systems to deliver engaging narrative experiences for entertainment, education, and training. However, there are a number of open research questions, including but not limited to the following:

- How and when to intervene in the virtual world on behalf of the user?
- How to generate narrative structures?
- How to encode human authorial intent and at what level of abstraction?
- How to incorporate believable characters into an interactive narrative framework?
- How to tailor narrative experiences?

Seeking the answer to these questions involves simultaneous pursuit of fundamental questions of artificial intelligence along with the art and creativity of developing compelling, playable demonstration systems.

The importance of the human author in the creation of interactive narratives is of critical concern to developers and researchers of interactive narrative systems; knowledge representation is not merely an implementation detail but an important aspect of the art required to create engaging interactive experiences. While branching story graphs and hypermedia interactive narratives have a low barrier to entry, interactive narratives utilizing intelligent experience manager agents have required substantive knowledge engineering expertise. To date interactive narrative systems involving intelligent systems are largely denizens of research labs, with one notable exception being the *Façade* interactive drama, publicly released in 2005. Can non-expert human authors create interactive narratives with intelligent experience management techniques? To shed light on this question is to determine whether the development of interactive narrative systems can become a mainstream form of creative expression. Promising directions toward this end attempt to ease the encoding of authorial intent using machine learning to train experience managers on narrative examples [12] or crowdsourcing narrative knowledge from the World Wide Web [15, 5].

Storytelling is an integral part of the human experience. We communicate through stories, but also use stories to entertain and educate. Immersive, interactive systems that can effectively engage us within narrative experiences but still allow us to exhibit agency over our experiences have the potential to revolutionize the ways in which computational systems are used to entertain, educate, and train humans. Interactive narrative can be a compelling application domain through which to explore the fundamental questions of computational narrative intelligence.

## References

- [1] Ruth Aylett. Narrative in virtual environments — towards emergent narrative. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium (Technical Report FS-99-01)*, pages 83–86. AAAI Press, Menlo Park, 1999.
- [2] Joe Bates. Virtual reality, art, and entertainment. *Presence: The Journal of Tele-operators and Virtual Environments*, 1(1):133–138, 1992.
- [3] Richard J. Gerrig. *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. Yale University Press, New Haven, 1993.
- [4] Pablo Gervás. Computational approaches to storytelling and creativity. *AI Magazine*, 30(3):49–62, 2009.
- [5] Boyang Li, Stephen Lee-Urban, and Mark O. Riedl. Toward autonomous crowd-powered creation of interactive narratives. In *Proceedings of the 5th Workshop on Intelligent Narrative Technologies*, 2012.
- [6] Brian Magerko. Evaluating preemptive story direction in the interactive drama architecture. *Journal of Game Development*, 2(3):25–52, 2005.
- [7] Michael Mateas and Andrew Stern. Integrating plot, character, and natural language processing in the interactive drama Façade. In *Proceedings of the 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 2003.
- [8] Janet H. Murray. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. MIT Press, Cambridge, MA, 1997.
- [9] Mark Nelson and Michael Mateas. Search-based drama management in the interactive fiction Anchorhead. In *Proceedings of the 1st Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2005.
- [10] Julie Porteous, Marc Cavazza, and Fred Charles. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Transactions on Intelligent Systems and Technology*, 1(2):1–21, 2010.
- [11] Mark O. Riedl, Andrew Stern, Don M. Dini, and Jason M. Alderman. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on System Science and Applications*, 3(1):23–42, 2008.
- [12] David L. Roberts, Mark Nelson, Charles Isbell, Michael Mateas, and Michael Littman. Targeting specific distributions of trajectories in MDPs. In *Proceedings of the 21st National Conference on Artificial Intelligence*, 2006.
- [13] Jonathan Rowe, Lucy Shores, Bradford Mott, and James Lester. Integrating learning, problem solving, and engagement in narrative-centered learning environments. *International Journal of Artificial Intelligence in Education*, 21:115–133, 2011.
- [14] Manu Sharma, Santi Ontañón, Manish Mehta, and Ashwin Ram. Drama management and player modeling for interactive fiction games. *Computational Intelligence*, 26(2):183–211, 2010.
- [15] Ried Swanson and Andrew Gordon. Say anything: A massively collaborative open domain story writing companion. In *First International Conference on Interactive Digital Storytelling*, 2008.

- [16] David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylishen. Interactive storytelling: A player modelling approach. In *Proceedings of the 3rd Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2007.
- [17] Kurt VanLehn. The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, 16, 2006.
- [18] Peter Weyhrauch. *Guiding Interactive Fiction*. PhD thesis, Carnegie Mellon University, 1997.
- [19] R. Michael Young, Mark O. Riedl, Mark Branly, Arnav Jhala, R.J. Martin, and C.J. Saretto. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1:51–70, 2004.
- [20] Hong Yu and Mark O. Riedl. A sequential recommendation approach for interactive personalized story generation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi Agent Systems*, 2012.

**Mark Riedl** is an Assistant Professor at the Georgia Institute of Technology (School of Interactive Computing). He received his Ph.D. in computer science from North Carolina State University in 2004. Mark’s research resides at the intersection of artificial intelligence, storytelling, and virtual worlds, focusing on how intelligent systems can autonomously create engaging experiences for users in virtual worlds. This research has implications for entertainment computing and virtual educational environments.

**Vadim Bulitko** is an Associate Professor at the University of Alberta (Department of Computing Science). He received his Ph.D. in computer science from the University of Illinois at Urbana-Champaign in 1999. Vadim is interested in building the strong artificial intelligence as well as understanding intelligence and cognition in humans and animals. He specializes in real-time heuristic search, AI in computer games including interactive narrative, and cognitive processes and models.