

From Linear Story Generation to Branching Story Graphs

Mark O. Riedl

Institute for Creative Technologies
University of Southern California
13274 Fiji Way, Marina Del Rey, CA 90292, USA
riedl@ict.usc.edu

R. Michael Young

Department of Computer Science
North Carolina State University
Raleigh, NC 27695, USA
young@csc.ncsu.edu

Abstract

Narrative is an important part of the way we interact with and make sense of the world. Interactive narrative systems tell stories in a virtual world in which the user is an interactive participant. Since the behaviors the user performs in the virtual world can affect the way in which a storyline unfolds, interactive narrative systems often use a branching story structure where non-interactive story presentations are interleaved with user decision points. An alternative approach – narrative mediation – represents story as a *linear* progression of events with anticipated user actions and system-controlled agent actions together in a partially-ordered plan. For every possible way the user can violate the story plan, an alternative story plan is generated. If narrative mediation is powerful enough to express the same interactive stories as systems that use branching story structures, then linear narrative generation techniques can be applied to interactive narrative generation with the use of narrative mediation. This paper sketches out a proof that narrative mediation is at least as powerful as acyclic branching story structures.

Introduction

Narrative as entertainment, in the form of oral, written, or visual stories, plays a central role in our social and leisure lives. Narrative is also used in education and training contexts to motivate and to illustrate. Blair and Meyer (1997) coined the term “narrative intelligence” to refer to the ability – human or computer – to organize experience into narrative. A computer system that takes a narrative approach to entertainment, education, or training will benefit from the ability to reason about narrative intelligence because the system can structure its narrative in ways that afford understanding by the user.

Recently, narrative intelligence has been applied to virtual worlds in order to create *interactive narrative systems*. An interactive narrative system is a virtual world in which a story unfolds and the user is considered a character in the story, able to interact with elements and other characters in the virtual world. The standard approach to incorporating storytelling into a computer system, however, is to script a story at design time. That is, the system designers determine ahead of time what the

story should be and hard-code the story into the system. For a computer system to use a scripted story means that the ability of the system to adapt to the user’s preferences and abilities is limited. The story scripted into a system may not completely engage the user’s interests or may be too challenging for the user to follow. Furthermore, if stories are scripted at design time, a system can only have a limited number of stories it can present to the user. In entertainment applications, a limited number of stories or a limited number of permutations of a single story in a computer game results limited replay value of that game. In educational and training applications, a limited number of stories or a limited number of permutations of a single story results in a system that cannot fully cater to the student’s needs and abilities.

The alternative approach is to generate stories either dynamically or on a per-session basis (one story per time the system is engaged). Narrative generation is a process that involves the selection of narrative content (the events that will be presented to an audience), ordering of narrative content and presentation through discourse of narrative content. A system that can generate stories is capable of adapting narrative to the user’s preferences and abilities, has expanded “replay value,” and is capable of interacting with the user in ways that were not initially envisioned by system designers.

Interactivity and Narrative

There are two fundamental types of narratives used in computer games and education and training applications: linear narrative and branching narrative. Linear narrative is a traditional form of narrative in which a sequence of events is narrated from beginning to ending without variation or possibility of a user altering the way in which the story unfolds or ends. Computer games typically employ linear narratives although the story structure is partitioned into interactive portions – levels – and cut-scenes. Even though the user has a certain degree of control during level play, the only outcome is successful completion of some objective (usually killing all the enemies in an area) or failure, in which case the user must try again. All users experience the same story and each user will experience the same story during successive sessions.

Some computer games use branching narrative in which there are many points in the story at where some action or

decision made by the user alters the way in which a narrative unfolds or ends. Branching narratives (e.g. Galyean 1995; Gordon et al. 2004) are typically represented as directed graphs in which each node represents a linear, scripted scene followed by a decision point. Arcs between nodes represent decisions that can be made by the user. Even though a branching narrative may introduce variability into the experience a user has with a storytelling system, the variability is scripted into the system at design time and is thus limited by the system designer's anticipation of the user's needs or preferences. The system pieces together pre-scripted sequences based on feedback from the user. The computational complexity of branching stories is such that story graphs scripted at design time have either a low branching factor or a limited number of decision points (Bruckman 1990). The user is constrained to the structure of the branching story graph constructed by the system designer such that users that make the same choices at each decision point will have identical experiences with the system. That is, if a user were to make the same decisions during two consecutive sessions with the system, her experience would be the same.

Interactive narrative systems not only have to consider the quality of the storytelling experience, but must balance the coherence of a story against the amount of control afforded the user (Riedl, Saretto, and Young 2003). The understandability of any narrative is determined in part by its *coherence* – the user's ability to comprehend the relationships between the events in the story, both within the story world (e.g., the causal or temporal relations between actions) and in the story's telling (e.g., the selection of camera sequences used to convey the action to the user). Systems that construct stories should respect the user's sense of coherence by clearly linking each action in the story world to its overall structure.

The degree of engagement by a user within an interactive narrative lies, to a great extent, with the user's perceived degree of *control* over her character as she operates within the environment. The greater the user's sense of control over her character, the greater will be her sense of presence (Lombard and Ditton 1997) – the sense that she is a part of the story world and free to pursue her own goals and desires.

Unfortunately, control and coherence are often in direct conflict in an interactive narrative system. To present a coherent narrative, the actions within an interactive narrative system's story are carefully structured (either at design time by human designers or at run time by narrative generation systems) so that actions at one point in the story lead clearly to state changes necessitated by actions occurring at subsequent points in the story. When users exercise a high degree of control within the environment, it is likely that their actions will change the state of the world in ways that may interfere with the causal dependencies between actions as intended within a storyline.

Narrative mediation (Riedl, Saretto, and Young 2003; Young et al. 2004) is a technique whereby linear narratives are made interactive. The question is: Is the expressive

power of narrative mediation at least as powerful as the story graph representation? If "yes," then linear narrative generation can be applied to the generation of branching narrative. The proofs that follow show that narrative mediation is at least as expressive as acyclic branching stories and also provides a feature that is not typically considered by branching stories: concurrency of user and character action.

Generating Interactive Narratives

Autonomous Agents

One technique for generating interactive narrative is to implement the system-controlled story-world characters as autonomous agents that are capable of reacting to the user and the environment in a believable manner. The story emerges from the decisions the autonomous agents make and the behaviors they perform in the virtual world (Aylett 2000). One advantage of this approach is that the characters in the story-world are capable of reacting to any actions the user performs. However, there is no explicit representation of plot or any defined notion of the outcome of the story; it is possible that no coherent narrative structure emerges.

In order to preserve the coherence of a system in which characters are implemented as autonomous agents, some interactive narrative systems (e.g. Kelso, Weyhrauch, and Bates 1993; Weyhrauch, 1997; Magerko et al. 2004) use *drama managers*. A drama manager (Kelso, Weyhrauch, and Bates 1993) is a process that monitors the activities of the autonomous characters and the user character relative to a partially ordered graph of events that move the story forward. When more than one event is possible, the drama manager analyzes the tradeoff of having one or the other next and subtly manipulates the world and autonomous characters so that the event that leads to the more satisfying experience occurs. A plot graph is not a representation of a branching story graph although adversarial-like search can be used to find complete linear narratives in the branching space of possible narratives (Weyhrauch 1997; Lamstein and Mateas 2004). Search of the complete narrative space is, unfortunately, intractable (Lamstein and Mateas 2004).

Narrative Mediation

Another technique for generating an interactive narrative is *narrative mediation* (Riedl, Saretto, and Young 2003; Young et al. 2004) which places control of character actions in the hands of a centralized "author" agent. The system generates a linear narrative that represents the ideal story that should be told to the user and then considers all the ways in which the interactive user can interact with the world and with the other characters. The generated story includes actions that system-controlled characters perform as well as actions that the user-controlled character should perform. For every action that the user makes that

threatens to deviate too severely from the linear story proposed by the system, the system dynamically generates an alternative storyline from the point of the deviation.

With narrative mediation, the story is represented by a plan. The plan contains annotations that explicitly mark the temporal relationships between all actions – user and system-controlled character – in the plan, defining a partial order indicating the steps’ order of execution. Other annotations, called causal links, are used to mark all causal relationships between the actions in the plan as well. A causal link connects two plan steps s_1 and s_2 via condition e , written $s_1 \xrightarrow{e} s_2$ when s_1 establishes the condition e in the world needed by subsequent action s_2 in order for s_2 to execute (Penberthy and Weld 1992). As the user issues commands for their character to perform actions in the story world, these actions must be checked against the story plan to determine whether they are *exceptions*. Exceptional actions have effects that threaten the conditions in the world required by future system-controlled character actions. Specifically, an exception occurs whenever a user attempts to perform some action α , where some effect $-e$ of α threatens to undo some causal link $s_1 \xrightarrow{e} s_2$ between two steps s_1 and s_2 , with condition e , where s_1 has occurred prior to α and s_2 has yet to occur.

Using planning structures to model narrative is advantageous because a narrative plan lays out the entire sequence of actions that will be performed during a storytelling session. The causal structure of the story plan is analyzed to determine all possible exceptions that can occur during the entire duration of the narrative. For every possible exception, an alternative story plan is generated that begins at the point of the exception. The process of narrative mediation defines a tree of partial story plans called a *narrative mediation tree* such that each plan represents a complete storyline, including both user actions and system-controlled character actions. The narrative mediation tree is constructed before execution of the interactive narrative session begins. To prevent the narrative mediation tree from growing infinitely large, some user actions are intervened with. Intervention is a process whereby a user action is surreptitiously replaced by a similar action, called a *failure-mode*, with different effects. The system maintains a list of failure-modes for each possible user action. For example, the user action Shoot has failure-modes Shoot-And-Miss and Gun-Jam. When intervention is determined to be the best policy for handling an exception, a failure-mode that has effects that do not threaten the plan is selected for execution instead. An exception that is intervened with does not necessarily require an alternative story plan since the causal structure of the original story plan is preserved.

Relationship between Narrative Mediation and Story Graphs

The technique of narrative mediation demonstrates that any system that can generate a linear narrative plan (with




Figure 1. A story graph.

causal annotation) coupled with a re-planning capability can be used to generate interactive narrative. Examples of linear narrative planners are Fabulist (Riedl and Young 2004; Riedl 2004) and Universe (Lebowitz 1985). However, many interactive narratives are expressed as story graphs; the use of branching story graphs is an intuitive way of expressing interactivity in storytelling. The question is: Is the expressive power of narrative mediation at least as powerful as the story graph representation? If the answer is “yes”, then narrative mediation can be used to generate any branching narrative structure that can be represented as a story graph. Such an answer would conclusively demonstrate that linear narrative generation can be applied to the generation of branching narrative.

A branching story structure is a story graph – a directed graph of nodes connected by arcs that represent user choices. Every possible path through the graph represents a story that can be told to the user. The user’s sense of control over the development of the story is limited by the number of arcs in a particular path in the branching story graph. An example of a story graph is shown in Figure 1. The system starts out non-interactively with system-controlled characters performing actions α_1 and α_2 . The user then chooses to perform action δ_1 or δ_2 . If δ_1 is chosen, system-controlled characters perform actions α_3 , α_4 , and α_5 . An ϵ -transition is taken in the absence of any user action.

Narrative mediation (Riedl, Saretto, and Young 2003; Young et al. 2004) is a technique for enhancing the user’s sense of control in an interactive narrative system. The idea behind narrative mediation is to generate a linear story structure that represents the best story that the system can tell. The linear story structure includes actions that the interactive user should perform. There is, however, the possibility that the user performs an exceptional action that interferes with the story structure. When an exception occurs, the system can either intervene or accommodate the action. Intervention means to prevent the exceptional action from interfering with the story structure. Accommodation means to incorporate the exceptional action into the story and generating a new linear story structure that is not threatened by the exception. Figure 2 shows an example of a narrative mediation tree. The node at the top of the graph represents the best linear story that the system can tell. The actions α_1 through α_6 are performed by both system-controlled characters and the user. Actions δ_1 , δ_2 , and δ_3 are exceptions. Note that there are many arcs for the δ_1 exception. This is because actions α_3 and α_4 in the node are unordered relative to each other;

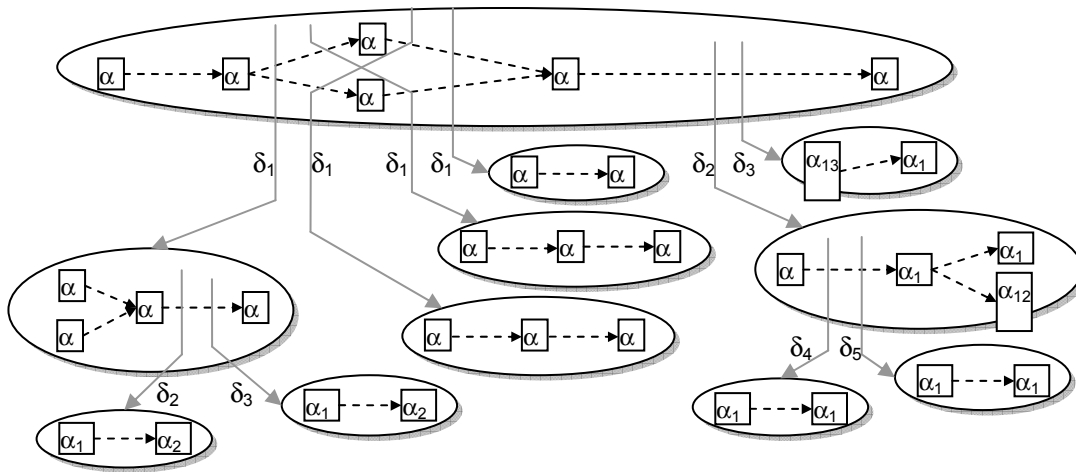


Figure 2. A narrative mediation tree.

which arc is followed and consequently which alternative story that is presented to the user depends on what actions have been executed prior to the exception. To provide the user with additional opportunities to exert control, the user can perform exceptional actions at any time, regardless of whether the linear story includes a user action or not.

Narrative mediation makes generation of interactive narratives possible by generating multiple contingent story structures. However, it is necessary to prove that narrative mediation is at least as expressive as a system that uses the more conventional story graph. For narrative mediation to be as expressive as a system that uses a story graph, for any possible story graph, a narrative mediation tree must exist that is equivalent. If narrative mediation is at least as expressive as story graphs, then a linear narrative generation system can be used to generate branching narratives through a process such as narrative mediation.

Definitions

The following are definitions necessary for the proofs. All graph structures used in the proof are assumed to use a basic partially-ordered plan structure to represent temporally ordered actions.

Definition 1 (Partially-ordered plan): A partially-ordered plan is a tuple $\langle A, O \rangle$ such that A is a set of actions performed in the story world and O is a set of ordering constraints of the form $\alpha_i < \alpha_j$ where $\alpha_i, \alpha_j \in A$.

A story graph G is a set of partially-ordered plans representing the non-interactive sequences of events. A set of partially-ordered plans $P \in G$ make up the nodes in the graph. The nodes are connected by story branches, which represent decision-points where the user can make a choice or perform some action. There are necessarily a finite number of actions $D \in G$, some of which are possible user choices after every partially-ordered plan executes.

Definition 2 (Story branch): A story branch in story graph G is a tuple $\langle p_1, C, \delta, p_2 \rangle$ such that $p_1, p_2 \in P$, $\delta \in D$, and C is an implementation-specific set of applicability criteria. The story branch indicates that p_1 is

temporally ordered before p_2 and that δ is an action executed by the user to cause transition from p_1 to p_2 .

Definition 3 (Story graph): A story graph is a tuple $\langle P, D, \Lambda, B \rangle$ such that P is a set of partially-ordered plans, D is a set of user actions, Λ is an action library such that for all $p_i \in P$, $A(p_i) \subseteq \Lambda$, and B is a set of story branches such that $C(b) = \emptyset$ for all $b \in B$.

The story graph representation does not disallow cycles. A cycle in a story graph means that there is a point in the story where the user can make a decision that causes a previously experienced portion of the story to repeat. An acyclic story graph, however, explicitly prohibits cycles, implying that a story has a finite duration. A story tree is a special type of acyclic story graph that does not reuse any node; each story branch terminates in a unique node.

The narrative mediation tree is the primary structure used by narrative mediation. While the data structure is superficially similar to a story graph, the story nodes are used in very different ways. Instead of being a short sequence of actions between decision points, a story node in narrative mediation is meant to represent the entire story if the user chooses not to interfere.

Definition 4 (Narrative mediation tree): A narrative mediation tree is a tuple $\langle P, D, \Lambda, B \rangle$ such that P is a set of partially-ordered plans, D is a set of user operations, Λ is an action library such that $D \subseteq \Lambda$ and for all $p_i \in P$, $A(p_i) \subseteq \Lambda$, and B is a set of story branches such that for all $b \in B$, $C(b) \subseteq A(p_i(b))$ and $C(b)$ is a prefix of $p_i(b)$.

The applicability criteria of b is the sequence of actions that form a prefix of the node from which b originates. $C(b)$ is the history of the actions that have been executed in the node that b originates from. The history is used to uniquely identify an arc in the case that several story branches with the same exceptional user action originate from the same node.

Proofs

The following are sketches of proofs that establish that narrative mediation trees are able to express at least the same set of interactive stories that acyclic story graphs can

Let G_a be an acyclic story graph.	
1. Let $G_i = \langle P, D, A, B \rangle$ be a story tree that is equivalent to G_a .	[Make a story tree out of the acyclic story graph.]
2. Let $D' = B' = \emptyset$. Let $B'' = B$. Let $P' = P$. Let $A' = A$.	
3. If $B'' = \emptyset$, then halt and return $G_m = \langle P', D', A', B' \rangle$. Otherwise, Let $b'' = \langle p_l, \emptyset, \delta, p_2 \rangle$ be an arc in B'' .	[Loop until B'' is empty. Otherwise, pick any story branch.]
4. Let $L \subseteq A(p_l)$ such that $\alpha_i \in L \rightarrow \langle \alpha_i, \alpha_j \rangle \notin O(p_l)$ for all $\alpha_i, \alpha_j \in A(p_l)$. Let $F \subseteq A(p_2)$ such that $\alpha_l \in F \rightarrow \langle \alpha_l, \alpha_i \rangle \notin O(p_2)$ for all $\alpha_l, \alpha_i \in A(p_2)$.	[Collect all “last steps” from the first node and all “first steps” from the second node in the branch.]
5. Let $p' = \langle A(p_l) \cup A(p_2) \cup \{\delta\}, O(p_l) \cup O(p_2) \cup \{\langle l, \delta \rangle \text{ for all } l \in L\} \cup \{\langle \delta, f \rangle \text{ for all } f \in F\} \rangle$. Let $A' = A' \cup \{\delta\}$.	[Create a new plan node that merges the two nodes.]
6. For all $b = \langle p_i, H, \delta_k, p_j \rangle \in B'$ such that $p_i = p_l$, $B' = B' - \{b\} \cup \{\langle p', H, \delta_k, p_j \rangle\}$.	[Update any story branches in mediation tree that originate from p_l .]
7. For all $b = \langle p_i, H, \delta_k, p_j \rangle \in B'$ such that $p_j = p_l$ or $p_j = p_2$, $B' = B' - \{b\} \cup \{\langle p_i, H, \delta_k, p' \rangle\}$.	[Update any story branches in mediation tree that terminate in p_l or p_2 .]
8. $B'' = B'' - \{b''\}$.	
9. For each $b = \langle p_i, \emptyset, \delta_k, p_j \rangle \in B''$ such that $p_i = p_l$, $B'' = B'' - \{b\}$, $B' = B' \cup \{\langle p', A(p_l), \delta_k, p_j \rangle\}$, and $D' = D' \cup \{\delta_k\}$.	[Other branches coming out of p_l become exceptions.]
10. For each $b = \langle p_i, \emptyset, \delta_k, p_j \rangle \in B''$ such that $p_i = p_2$, $B'' = B'' - \{b\} \cup \{\langle p', \emptyset, \delta_k, p_j \rangle\}$.	[Update any story branches in story tree that originate from p_2 .]
11. $P' = P' - \{p_l, p_2\} \cup \{p'\}$.	[Add new node to mediation tree.]
12. Go to step 3.	

Figure 3. Algorithm transforming an arbitrary acyclic story graph into a narrative mediation tree.

express. If the expressive power of narrative mediation trees is at least as powerful as that of acyclic story graphs, then narrative mediation can generate any interactive branching storyline that is used by a system that implements acyclic story graphs. In order to prove this, we show that the set of all acyclic story graphs is a subset of the set of all narrative-mediation trees. The subset relationship is true if there is an algorithm that transforms any arbitrary acyclic story graph into an equivalent narrative mediation tree representation. For an acyclic story graph to be equivalent to a narrative mediation tree, the set of all paths – the set of all stories that can be told – through a story graph must be a subset of the set of all paths through a narrative mediation tree. A path through either structure includes the system-controlled character actions interleaved with user actions. For example $\{\alpha_1, \alpha_2, \delta_2, \alpha_6, \alpha_7, \delta_5, \alpha_8, \alpha_9, \delta_3, \alpha_{10}, \alpha_{11}\}$ is one path through the story graph in Figure 1.

The following facts can be proven through simple proofs. Constraints on the length of this paper prohibit their inclusion here.

- **The set of all acyclic story graphs is a subset of all story graphs.**
- **The set of all acyclic story graphs is equal to the set of all story trees.** An acyclic graph can be transformed into a tree by duplicating nodes with incoming arcs.
- **The set of all narrative mediation trees is not a subset of the set of all story graphs.** Narrative mediation allows user actions to execute concurrently with system-controlled character actions while story graphs require user actions and system actions to be interleaved.
- **The set of all story graphs is not a subset of the set of all narrative mediation trees. Nor is the set**

of all narrative mediation trees a subset of the set of all acyclic story graphs. Story graphs can have loops and consequently have infinite-length paths while narrative mediation does not allow infinite stories.¹

The key to proving the hypothesis, however, is proving that the set of all acyclic story graphs is a subset of the set of all narrative mediation trees. To prove this, we provide an algorithm that can transform any arbitrary instance of an acyclic story graph into an equivalent narrative mediation tree.

Sketch of proof: The set of all acyclic story graphs is a subset of the set of all narrative mediation trees.

Let G_a be an acyclic story graph. The algorithm is shown in Figure 3 and produces a narrative mediation tree, G_m .

The proof of correctness of the algorithm in Figure 3 relies on the fact that a node in G_m contains both system-controlled character actions and user actions, representing one path through G_a . Line 5 creates a new node that is a concatenation of two nodes in G_a plus the user action that transitions between those two nodes. Line 11 removes the old nodes from G_m and adds the new amalgamated node. As a result, G_m has a minimal number of nodes that represent the maximum-length sub-paths in G_a .

Whether or not G_m is equivalent to G_a relies on the positioning of the arcs in G_m . If there are two arcs originating in a single node in G_a , then in G_m , the user action of one of the arcs must be part of a partially-ordered

¹ Cyclic story paths result in infinite-depth branches in a narrative mediation tree. In practice, we do not anticipate that users will attempt to loop forever, justifying limitation of tree depth. Interventions can be used to enforce depth limits.




Figure 4. Relationship between story graphs and narrative mediation trees.

plan in a story node in G_m while the other user actions are represented by arcs originating in that node in G_m . Line 5 incorporates one user action into a new node. Line 9 identifies the alternative arcs in G_a (because the arc originates from the same node in G_a) and creates a new arc in G_m originating from the new node.

The proof of correctness also relies on the correct determination of the applicability constraints (Line 9) for all story branches in the resulting narrative mediation tree. The correctness can be proven directly, but is beyond the scope of this paper.

Consequently G_a and G_m are equivalent. Therefore, any arbitrary acyclic story graph can be transformed into an equivalent narrative mediation tree. Therefore the set of all acyclic story graphs is a subset of all narrative mediation trees. □

Conclusions

The conclusion that can be drawn from the proofs in the previous section is that narrative mediation is at least as powerful as interactive narrative systems that have acyclic branching stories. Any acyclic story graph can be transformed into an equivalent narrative mediation tree. Narrative mediation can represent stories in which user actions are performed concurrently with system-controlled character actions, something that is impossible to represent in the standard story graph representation. In that respect, narrative mediation trees can represent interactive stories that cannot be represented by story graphs. However, narrative mediation trees, by nature, cannot have cycles, meaning there is a class of story graph – cyclic story graphs – that cannot be represented by narrative mediation trees. The relationship between story graphs and narrative mediation trees is shown in Figure 4.

Narrative mediation uses a narrative generator to construct linear narratives that are organized as a tree of contingencies. Since there is an equivalent narrative mediation tree for every acyclic story graph, interactive stories can be generated by iteratively invoking a linear narrative generator such as that used in Fabulist (Riedl and Young 2004; Riedl 2004) that represents the causal relationships between story world events.

Acknowledgements

The work presented here has been supported by the National Science Foundation CAREER Award #0092586.

Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government and no official endorsements should be inferred.

References

- Aylett, R. 2000. Emergent narrative, social immersion, and “storification.” *Proc. of the 1st Int. Workshop on Narrative and Interactive Learning Environments*.
- Blair, D. & Meyer, T. 1997. Tools for an interactive virtual cinema. In R. Trappl & P. Petta (Eds.) *Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents* (pp. 83-91). New York: Springer.
- Bruckman, A. 1990. *The Combinatorics of Storytelling: Mystery Train Interactive*. Unpublished manuscript.
- Galyean, T.A. 1995. *Narrative Guidance of Interactivity*. Ph.D. dissertation, MIT Media Arts and Sciences Program.
- Gordon, A., van Lent, M., van Velsen, M., Carpenter, P., & Jhala, A. 2004. Branching storylines in virtual reality environments for leadership development. *Proceedings of the 16th Innovative Applications of AI Conf.*
- Kelso, M.T., Weyhrauch, P., & Bates, J. 1993. Dramatic presence. *Presence: The Journal of Teleoperators and Virtual Environments*, 2(1).
- Lamstein, A. & Mateas, M. 2004. Search-based drama management. *Proc. of the 2004 AAAI Workshop on Challenges in Game Artificial Intelligence*.
- Lebowitz, M. 1985. Story-telling as planning and learning. *Poetics*, 14, 483-502.
- Lombard, M. & Ditton, T. 1997. At the heart of it all: The concept of presence. *Journal of Computer-Mediated Communication*, 3(2).
- Magerko, B., Laird, J.E., Assanie, M., Kerfoot, A., & Stokes, D. 2004. AI characters and directors for interactive computer games. *Proc. of the 16th Innovative Applications of AI Conference*.
- Penberthy, J.S. & Weld, D. 1992. UCPOP: A sound, complete, partial-order planner for ADL. *Proc. of the 3rd Int. Conf. on Knowledge Representation and Reasoning*.
- Riedl, M.O. 2004. *Narrative Planning: Balancing Plot and Character*. Ph.D. dissertation, North Carolina State University.
- Riedl, M.O., Saretto, C.J., & Young, R.M. 2003. Managing interaction between users and agents in a multiagent storytelling environment. *Proc. of the 2nd Int. Joint Conf. on Autonomous Agents and Multi Agent Systems*.
- Riedl, M.O. & Young, R.M. 2004. An intent-driven planner for multi-agent story generation. *Proc. of the 3rd Int. Joint Conference on Autonomous Agents and Multi Agent Systems*.
- Weyhrauch, P. 1997. *Guiding Interactive Fiction*. Ph.D. dissertation, Carnegie Mellon University.
- Young, R.M., Riedl, M.O., Branly, M., Jhala, A.H., Martin, R.J., & Saretto, C.J. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1(1).