

Automated Story Direction and Intelligent Tutoring: Towards a Unifying Architecture

Mark RIEDL, H. Chad LANE, Randall HILL, William SWARTOUT
Institute for Creative Technologies, University of Southern California
13274 Fiji Way, Marina Del Rey, CA 90292, USA

Abstract. Recently, interactive storytelling systems – systems that allow a user to make decisions that can potentially impact the direction of a narrative – have been applied to training and education. Interactive storytelling systems often rely on an automated story director to manage the user’s experience. The focus of an automated director is the emergence of a narrative-like experience for the user. In contrast, intelligent tutors traditionally address the acquisition or strengthening of a learner’s knowledge. Our goal is to build training simulations that cultivate compelling storylines while simultaneously maintaining a pedagogical presence by incorporating both automated story direction and intelligent tutoring into an immersive environment. But what is the relationship between an automated director and an intelligent tutor? In this paper, we discuss the similarities and differences of automated story directors and intelligent tutors and, based on our analysis, recommend an architecture for building narrative-based training simulations that utilize both effectively and without conflict.

Introduction

Narrative, in the form of oral, written, or visual stories, plays a central role in our social and leisure lives. Although stories are often created with the goal of entertainment, many also *happen* to be educational because of their content. Another class of narratives, the kind we focus on, are those that are *intended* to be educational. In these cases the inherent benefits of narrative, such as increased motivation and improved context for illustration, are leveraged for achievement of the learning objectives. Narrative-centered learning environments [7] have the potential to go beyond traditional, non-interactive narrative forms by enabling a student to be an interactive part of a story. In a narrative-centered learning environment, the purpose of story is to engage the student in a series of challenging problem-solving activities and/or dilemmas. Each of these hurdles is intended to have some educational value. However, when pedagogy becomes entwined with narrative, it becomes unclear as to whether a narrative-centered learning environment is performing storytelling, tutoring, or a little of both. Research in interactive storytelling systems (e.g. [18; 10; 15; 19; 8]) and intelligent tutoring systems [17; 11] has remained largely separate. Recently, training and educational systems such as those described in [9], [14], [12], and [5] involve elements of both.

We are interested in building narrative-based training simulations that accelerate students’ problem solving skills in complex, real world situations where quick and successful resolution of problems requires the application of tacit knowledge. For example, training for military leaders deployed into regions where they must interact and negotiate with people of different cultures. Storytelling is an effective means for transferring tacit knowledge [13]. More traditional training techniques involve simulators that run fixed scenarios. Simulators are effective practice tools because they enable a trainee to practice skills, try different strategies, and discover outcomes. However, it can also be advantageous for trainees to be situated in an environment that promotes a more targeted and reproducible experience. The

benefit of a narrative-based approach is that the trainee can be exposed to a larger context with multiple learning objectives that are strung together for greater pedagogical impact, providing valuable training experience dealing with life-like successions of interrelated problems. For example, a trainee negotiates with tribal leaders to achieve some effect and then transitions to a seemingly unrelated problem with security lapses only to discover that the side-effects of the first situation impact the way the second situation is resolved. For a simulation to evolve in a globally structured way, some supervisory control is required to coordinate the activities of story participants, such as a “director” of a movie. Specifically, the purpose of an automated story director is to coerce a virtual environment into producing certain events or achieving a certain outcome. The focus of the automated director is on the *experience* of the trainee, not the acquisition or strengthening of knowledge. In contrast, the acquisition and strengthening of the trainee’s knowledge is precisely the role intelligent tutoring systems traditionally play in learning systems. But what is the relationship between an automated story director and an intelligent tutor?

We believe that narrative-based training simulations can benefit from having both an automated story director and an intelligent tutor. The purpose of this paper is to explore the similarities, differences, and relationship possibilities between automated story directors and intelligent tutors as they apply to narrative-based training simulations. As a result of our analysis presented below, we believe there must be a clear delineation of responsibilities and an understanding of how each can impact the other (either directly or indirectly). We present one possible delineation of responsibilities for automated story directors and intelligent tutors in a narrative-based training simulation. Finally, we propose an architecture for incorporating an automated story director and intelligent tutor into a narrative-based training simulation that implements our recommendations.

1. Interactive Narrative

In systems that employ storytelling techniques, there are two fundamental types of narratives used: 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. For a narrative to be interactive, it is necessary to introduce the notion of branching. This is accomplished by introducing decision points into the story such that user actions in the environment can potentially alter the sequence of events that is narrated.

An interactive narrative system is a computer program that manages a branching narrative and provides a facility for a user to influence the story. Typically, branching narrative systems achieve interactivity by situating the user as an embodied character in the story world. There is a distinction, however, between an interactive narrative and a simulation. From the perspective of the user, both simulations and interactive narratives involve the emergence of narrative-like situations. However, most interactive narrative systems employ an automated story director – also called a drama manager [6] – that manages the interactive virtual environment in order to coerce the emergence of narrative to correlate with some path through a given model of branching narrative. The automated story director has a global perspective of the story world which it attempts to reconcile with some representation of branching narrative. There are many approaches to interactive narrative including: traversing a pre-scripted directed graph of scenes and choice-points (e.g. choose-your-own-adventure stories); dynamically generating alternative story branches (e.g. [19]); comparing emergent story to partially ordered event graphs (e.g. [18; 8]); modelling dramatic tension [15]; and dynamically choosing from a pool of scenes (or scene-like constructs such as beats) based on analysis of user interaction (e.g. [10]).

2. Intelligent Tutoring Systems

Intelligent tutoring systems (ITS) are usually defined as systems that use the techniques of artificial intelligence to evaluate student performance and offer suggestions for improvement. ITS has an extensive history spanning about four decades, and a huge number of systems have been developed during this time [17; 11]. The content and form of a tutoring session depend on many factors, such as when tutoring occurs (e.g. before, during, or after an exercise), level of guidance (some exploration vs. tighter control), and the kind of desired interactions (e.g., dialogue vs. graphical gestures), to name just a few. For this paper, we will mostly focus on one common form of tutoring known as *coached problem solving* [16]. Under this definition, the tutor and student collaborate to solve a problem. During this process, initiative shifts back and forth: as long as the student is taking correct steps, the tutor simply indicates agreement or remains silent. If the student becomes stuck or requests help, the tutor provides hints to get the student back on a correct solution path.

An ITS technology related to coached problem solving is *model tracing*, which attempts to understand and track students' actions at the cognitive level [1]. Model tracing tutors require an executable cognitive model of the target skill that is able to solve problems in the domain along with help messages that are displayed when necessary (e.g., when the student is stuck, or after a mistake). A common (but not universal) trait of model tracing systems is *immediate feedback*. When a student makes an error, the system will quickly let the student know about the mistake and help them repair it. Generally, the student is not allowed to continue without fixing the error. To provide this kind of help, model-tracing tutors maintain a step-by-step understanding of where the student is during problem solving. The goal of the tutor is to keep the student on a known solution path much or all of the time.

3. Distinctions between an Automated Story Director and Intelligent Tutor

As mentioned, building a narrative environment for learning naturally raises the question about where to draw the line between an automated director and tutor. There are many apparent similarities and differences between their roles. In this section, we lay out this space of functionality and requirements to act as a backdrop for our proposed architecture.

3.1. Tracking the user

An obvious similarity between automated directors and intelligent tutors is that both tend to maintain representations of valid “paths” through some space, and attempt to track the user as she performs actions in the environment. In the case of a tutor, this is usually a path through a solution space for a given problem – arcs correspond to problem solving steps and nodes to states of the evolving solution. The target knowledge consists of the individual rules for updating the state and the knowledge of how to apply them productively to achieve a *goal state*. The goal of the tutor is to impart its model of expert performance. Acquisition of this knowledge typically occurs through repetitive practice, and thus many ITSs present the student with many problems to solve.

In the case of an automated story director, the representation of valid paths can be a branching narrative as a path through state space of partially ordered steps that constrain the possible narratives. The situation is conceptually similar to that of an intelligent tutor's in that the user moves from state to state by taking actions within the environment. For example, the user may choose to move from one room to another in search of some desired item. However, the actions taken by the user are *not* part of some underlying expert model of performance. The purpose of a director is not to communicate *how* to participate in a

story, but rather to ensure that the story remains meaningful without violating the user's sense of control.

3.2. Overt versus Covert Intervention

Perhaps the greatest distinction between automated story directors and intelligent tutors is the nature in which they intervene and provide feedback to the user. Intelligent tutors, in a way similar to the human counterparts, maintain a very clear, overt pedagogical presence. The purpose of a tutor is to provide scaffolding to the student to reinforce the knowledge and concepts required to perform a skill. Tutors intervene when students make mistakes or deviate from known solution paths, although many systems delay this feedback to give the student a chance to self-correct. Human tutors can also provide positive feedback to reinforce correct actions, sometimes through subtle discourse acts, nodding, or even facial expressions.

Automated story directors, on the other hand, influence the emergence of story through *only* covert intervention. That is, a story director subtly manipulates the user through adjustments to the story world environment and through adjustments to the goals, beliefs, and desires of non-player characters (NPCs) in order to encourage the emergence of a story that corresponds to the story director's model of plot development [6]. An automated director is considered successful if the user's experience reaches a desirable outcome and the user is never aware of the director's presence.

3.3. Correcting versus Pushing

How the automated director and tutor handle situations where the user deviates from an anticipated path is also important. An intelligent tutor is concerned more with the performance of the user and attempts to guide through direct intervention. When the student makes a mistake, the tutor may ask a question, provide a hint, enter into a remedial dialogue with the student, or simply give the correct action away. This sometimes involves restoring the solution to some previous state. The tutor has control over the problem-solving environment and this kind of intervention does not violate the student's expectations in any way.

A story director, on the other hand, is concerned with the experience of the user and attempts to guide through subtle manipulation of the environment. Story directors do not intervene or require "undo" actions by the user since this would violate the user's sense of immersion. User actions that do not correspond to any branch in the narrative model are not necessarily problematic. Instead, a story director may push the user towards a world state that facilitates story progression – possibly one that is nearly identical to a previous world state. The subtle difference is that the previous actions remain part of the history and therefore part of the emergent story even while the user is covertly manipulated towards achieving some more desirable world state.

3.4. Scope of Control

For most ITSs the problem-solving environment exists solely to serve the purpose of instruction. It is therefore entirely reasonable for a tutor to intervene and even undo a student's action to give her another chance because there is no aspect of the world that is not controlled by the student or the tutor or both. There are, however, learning domains such as leadership training that involve problems that exist within a larger context of a real-world-like simulation. In such a domain, being immersed in the real-world-like simulation is just as important as solving the problems that occur within the larger context.

In such a domain, there are elements in the world that are not under the control of the student or the tutor such as non-player characters, weather, etc. In order to maintain the trainee's sense of immersion in the environment and in the role she is playing, a certain degree of believability is necessary. Specifically, the trainee must believe that the world is not a "toy" world, meaning the tutor will not encourage or allow "undo" actions when the trainee makes a mistake. The tutor has the ability to influence the trainee's problem solving, such as offering private consultation during a negotiation. The automated story director is, however, well suited to exert control over many aspects of the virtual environment. The story director must be limited so that it cannot influence the world in such a way that trivializes learning objectives, even if the trainee is required to complete a certain task for the plot to progress as the story director desires. That is, the scope of control of the automated director must be limited to only helping establish the conditions in which training objectives become relevant.

4. Building Narrative-Based Training Simulations

We are interested in narrative-based training simulations in which the trainee has a story-like experience. The purpose of such a system is to situate the trainee in a virtual world similar to that in which the trainee will be expected to operate in the future. During that experience, the trainee will be required to perform procedures that meet certain training objectives. The narrative aspect provides a layer of control over the simulation so that certain situations are guaranteed to arise and challenge the trainee. A narrative-based training simulation therefore requires an automated director. Although such a system would have pedagogical value by itself, we believe it is also beneficial to provide an intelligent tutor to help the trainee when necessary and improve the chances that the intended learning outcomes are realized.

4.1. Delineation of Story Director and Intelligent Tutor Responsibilities

Given the similarities between automated story directors and intelligent tutors and the possibility of overlap, we propose the following delineation of responsibilities. The purpose of the story director is to ensure the trainee has a valuable experience that places the trainee in certain pedagogically valuable situations (although it is not the responsibility of the automated director to assess the pedagogical value of the experiences themselves). The story director will achieve this by performing two responsibilities:

1. Representing the progression of the user's experience as a high-level narrative structure that is either linear or branching.
2. Guiding semi-autonomous agents so that they facilitate plot development as well as constraining them so that they cannot perform behaviors that directly conflict with the narrative structure of the automated director.

The story director maintains a representation of the desired plot. Aylett [2] recommends that plot occur at a high level and that everything else below is handled reactively. Here, "high-level story representation" means that the story representation does not contain instantiations of user-performed or non-player character performed operations that are also part of a learning task procedure. A non-player character (NPC) is a character in the virtual world that is not controlled by a human user. Suppose the story director requires a situation in which the trainee must convince an NPC of some fact in order to achieve some effect on the world. The plot representation can include a plot point where an NPC is convinced of the fact but should not include explicit representations of a specific set of low-level behaviors such as appeasement and negotiation that achieve that world state. That way, the story director does not require the trainee to perform a learning task in a particular way, constraining the functioning of the intelligent tutor. It also prevents the story director from pushing the virtual

environment in such a way that trivializes a learning objective. For example, the automated director could unknowingly interfere with a pedagogical goal by making NPCs agreeable when the student needs more of a challenge. The high-level representation also makes the automated director resilient to the trainee's fumbling of training procedures. For example, if the trainee struggles excessively during some task, it does not necessarily constitute a deviation from the story director's plot representation as long as the state of the world is such that future high-level plot points remain achievable.

The principle role of the automated director is to coerce certain situations – those represented in the high level narrative model – to develop that ordinarily would not develop in the context of an unguided simulation. This is best achieved through direction of the NPCs. We believe that NPCs should be controlled by autonomous agent technologies in order to be believable – to appear motivated by their own internal goals, desires, beliefs, personalities, etc. Fully autonomous agents, however, may not act in accordance with the story director's plot model and may even cause exceptions. The primary means for the story director to manipulate the world according to its representation of plot is through the NPCs. The automated director directs the agents from time to time to adopt goals that establish world states that facilitate plot advancement. For example, if the plot representation demands that a situation develops where a boy is hurt and needs medical assistance, the story director can insist that the agent controlling the boy NPC achieve the goal of becoming hurt near the trainee. Just as important, the director constrains the agents so that they cannot perform behaviors that directly contradict the narrative. For example, in order to force the trainee to deal with the situation with a boy that becomes hurt nearby, the automated director can require the mother of the boy NPC not seek immediate medical attention even though that may be the inclination of the agent controlling the mother NPC.

Intelligent tutoring can play at least two roles in a narrative-based training simulation. First, it is beneficial to have an intelligent tutor agent acting as an expert that accompanies the trainee to provide feedback and guidance. This tutor could be embodied as a character in the world, or merely a “disembodied voice in the trainee's ear.” The distinction of the tutor as an expert is important because it should be able to handle situations as they arise and provide scaffolding for the trainee as she attempts to work through them. The second role intelligent tutoring plays in a narrative-based training simulation is as a reflective tutor during an after action review (AAR). AARs allow the trainee to reflect on (with the help of the tutor) what happened, the quality of the student's performance, and how to improve in the future. With access to the automated director's logs, the tutor can also provide “behind the scenes” clarification on how the plot developed and give explanations regarding the causal relationships between events. This opens up the possibility of discussing 2nd and 3rd order effects of the user's actions as well as actions by NPCs outside the local perspective of the trainee that had a direct causal influence on the situations that the trainee encountered.

4.2. An Architecture for Narrative-Based Training Simulations

An architecture for a narrative-based training simulation that reflects the design decisions and delineation of responsibilities of a story director and intelligent tutors discussed above appears in Figure 1. It is based on four main components: the world simulation, the automated director, the intelligent tutor, and a collection of autonomous agents. The dashed links are optional. The world simulation implements the physics of the virtual world, graphical rendering, etc. Within the world simulation are models of the dynamic objects in the world. Most important of these objects are the user avatar, the non-player character (NPC) avatars – also called “bots” – and the intelligent tutor's world representation (either an avatar or a “disembodied voice”). The user controls her avatar through a user interface and receives a visualization of the virtual




Figure 1. The narrative-based training simulation architecture.

world in first-person or limited third-person perspective. The intelligent tutor’s world representation can be a fully embodied avatar in the case that the intelligent tutor also plays the role of a character or can be a disembodied presence that can “speak into the user’s ear,” modify the user’s visual interface, etc. The NPC bots are embodiments of characters that are controlled by the external AI autonomous agents.

The automated story director maintains a branching narrative model and monitors the world simulation relative to this model. We believe that the automated director should not interfere with the physics of the world simulation beyond extra-diegetic factors – elements that involve dramatic presentation such as music, lighting, and camera position. In addition to tracking the trainee in the world simulation, the automated director is also responsible for managing the NPCs. The NPCs are controlled by autonomous agents that are believable [3] and directable. A directable agent is one who’s reasoning and behavior can be specified or modified by an external process [4]. How to implement directable, semi-autonomous agents is beyond the scope of this paper.

The intelligent tutor, during the exercise, models the user’s knowledge structures and learning as they pertain to particular learning objectives. These learning objectives occur throughout the narrative experience although they are not strictly represented as learning objectives in the automated director’s narrative representation. The world simulation feeds the tutor relevant information, including learning task start and stop times, details of the environment and NPC actions during the task, and user actions (for evaluation and tracking). When the tutor needs to intervene (to provide solicited or unsolicited feedback), it does so through the world simulation.

The disambiguation of the story director’s and intelligent tutor’s roles in a narrative-based training simulation means that the story director and intelligent tutor do not need to communicate directly with each other. However, we recognize that it can be beneficial for information to flow between the two components. The intelligent tutor can be more effective if it receives projections of upcoming narrative events so that it can tailor any assistance it provides the user, as long as it doesn’t act on this omniscience to interfere with the plot progression. The automated director can also accept suggestions from the tutor to modify the narrative structure to be more pedagogically relevant to the trainee.

5. Conclusions

A narrative-based training simulation can benefit from the inclusions of both an automated story director and an intelligent tutor. Despite the outward appearance of overlap between narrative experience control and knowledge scaffolding and strengthening, there are enough differences to suggest that automated story directing and intelligent tutoring remain distinct in systems that use story to compel and motivate learning objectives. We recommend that a narrative-based training simulation incorporate both automated story direction and intelligent tutoring. We have laid out a proposal for the delineation of responsibilities that facilitates the inclusion of story control and tutoring side-by-side in a narrative-based training simulation.

Acknowledgements

The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

- [1] Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). *Cognitive tutors: Lessons learned*. The Journal of the Learning Sciences, 4 (2) 167-207.
- [2] Aylett, R. (2000). Emergent narrative, social immersion and “storification.” *Proceedings of the 1st International Workshop on Narrative and Interactive Learning Environments*.
- [3] Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7), 122-125.
- [4] Blumberg, B.M., & Galyean, T.A. (1995). Multi-level direction of autonomous creatures for real-time virtual environments. *Proceedings of SIGGRAPH'95*.
- [5] Johnson, W.L., Vilhjálmsson, H., and Marsella, S. (2005). Serious games for language learning: How much game, how much AI? *Proc. of the 12th Int. Conf. on AI in Education*.
- [6] Kelso, M.T., Weyhrauch, P., & Bates, J. (1993). Dramatic presence. *Presence: The Journal of Teleoperators and Virtual Environments*, 2(1): 1-15.
- [7] Mott, B.W., Callaway, C.B., Zettlemoyer, L.S., Lee, S.Y., & Lester, J.C. (1999). Towards narrative-centered learning environments. *Proceedings of the 1999 AAAI Fall Symposium on Narrative Intelligence*.
- [8] Magerko, B., Laird, J., Assanie, M., Kerfoot, A. & Stokes, D. (2004). AI characters and directors for interactive computer games. *Proc. of the 16th Innovative Applications of AI Conference*.
- [9] Marsella, S., Johnson, W.L., and LaBore, C. (2000). Interactive pedagogical drama. *Proc. of the 4th Int. Conf. on Autonomous Agents*.
- [10] Mateas, M., & Stern, A. (2003). Integrating plot, character, and natural language processing in the interactive drama Façade. *Proc. of the 1st Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment*.
- [11] Shute, V. J. and Psotka, J. (1996). Intelligent Tutoring Systems: Past, Present and Future. In D. Jonassen (ed.) *Handbook of Research on Educational Communications and Technology*. Scholastic Publications.
- [12] Sobral, D., Machado, I., & Paiva, A. (2003). Machiavellian characters and the edutainment paradox. *Proc. of the 4th Int. Working Conf. on Intelligent Virtual Agents*.
- [13] Sternberg, R., Forsythe, G., Hedlund, J., Horvath, J., Wagner, R., Williams, W., Snook, S., & Grigorenko, E. (2000). *Practical Intelligence in Everyday Life*. New York: Cambridge University Press.
- [14] Swartout, W., Hill, R., Gratch, J., Johnson, W.L., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Rickel, J., Thiébaux, M., Tuch, L., Whitney, R., and Douglas, J. (2001). Toward the holodeck: Integrating graphics, sound, character, and story. *Proc. of the 5th Int. Conf. on Autonomous Agents*.
- [15] Szilas, N. (2003). IDtension: A narrative engine for interactive drama. *Proc. of the 1st Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment*.
- [16] vanLehn, K. (1996). Conceptual and meta learning during coached problem solving. *Proc. of the 3rd Int. Conf. on Intelligent Tutoring Systems*.
- [17] Wenger, E. (1987) *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. San Francisco: Morgan Kaufmann.
- [18] Weyhrauch, P. (1997). *Guiding Interactive Fiction*. Ph.D. Dissertation, Carnegie Mellon University.
- [19] Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., & Saretto, C.J. (2004). An architecture for integrating plan-based behaviour with interactive game environments. *Journal of Game Development*, 1(1): 51-70.