Maintaining Cognitive Engagement in Training Scenarios Using Explicit Cognitive Models

Eric Roberts

Institute for Defense Analyses
Alexandria, VA
Eric.Roberts@adlnet.gov

R. Michael Young

North Carolina State University
Raleigh, NC
Young@csc.ncsu.edu

ABSTRACT

Commonly, military training games and simulations depend on participants engaging in the immersive environment and then stopping to glean the meaning of their behaviors through (admittedly, increasingly sophisticated) After Action Reports. In order to derive the meaning of their experience, they must break cognitive engagement with their experience. This is, of course, sub-optimal. While it is established that reflection-upon-behavior is required to "make sense" of experience, it may not be possible to demonstrate how that can be realized without breaking the participants' cognitive engagement.

This presentation demonstrates a unique method of maintaining user engagement through a planned system of "graceful failures," that allow non-catastrophic mistakes, precludes catastrophic mistake, and maintains "play" of the simulation.

Plan-based models of narrative control interaction within a learning environment to provide powerful underpinnings for models of both the environmental dynamics and the cognitive model of the learner operating within it. The idea of narrative mediation – an analysis of the potential points of failure within an automatically generated learning experience that pre-computes appropriate story adaptations at points where user activity could cause a learning experience to break.

This approach tracks every learner behavior in the engagement, making it possible to distinguish between the behavior-as-behaved and the behavior-as-instructionally-significant. This probable discrepancy typically is not addressed in games and simulations, leading to the real risk that active engagement in the game-play aspect is misconstrued as engagement in lessons to be learned.

ABOUT THE AUTHORS

R. Michael Young is an associate professor of Computer Science at North Carolina State University where he is co-director of the NCSU Digital Games Research Center and directs the Liquid Narrative research group. Michael is the Editor-in-Chief of the Journal of Game Development, was a founder and Conference Chair for the First Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE) in 2005 and served as the Tutorials Chair for AIIDE-2006. In 2007, he served as the Program Chair for the Third Annual Conference on Game Development in Computer Science Education and will serve as the Program Chair for the Fourth International Conference on the Foundations of Digital Games in 2009. Young has published a range of scientific and technical papers in the areas of interactive narrative, automatic 3D camera control, planning and computational linguistics. He was awarded a National Science Foundation CAREER Award in 2000.

Eric Roberts, Ph. D. is Chief Scientist for Learning, at the Advanced Distributed Learning Initiative, in Alexandria, Virginia. He is concerned with the preservation of instructional quality

when programs are moved from traditional design approaches to approaches conformant to SCORM and the nature of narrative to establish meaning.

Maintaining Cognitive Engagement in Training Scenarios Using Explicit Cognitive Models

Eric Roberts

Institute for Defense Analyses
Alexandria, VA
Eric.Roberts@adlnet.gov

R. Michael Young

North Carolina State University
Raleigh, NC
Young@csc.ncsu.edu

EXPERIENTIAL EDUCATION IS KEY TO MILITARY PREPAREDNESS

Knowing "how" and knowing "why", beyond knowing "what" to do is fundamental to military training and education – as it is for any practitioner faced with un-anticipated challenges. A study of classroom school teachers, for example (Kennedy, 1987) found that people who know technique but not theory are incapable of improvising.

As long ago as the 1920's John Dewey argued that one should "learn by doing." Recent studies of individuals using game-based learning suggest, however, that people really need, also, to learn by thinking about what it is they are doing (Squire, 2006).

In a thorough 1999 ARI report, Morrison and Meliza document the essential role of the After Action Report (AAR) in this light (Morrison and Meliza, 1999).

"As defined in Training Circular (TC) 25-20, A Leader's Guide to After-Action Reviews, (U.S. Army Combined Arms Center [CAC], 1993), an AAR " . . . is a professional discussion of an event, focused on performance standards, that enables soldiers to discover for themselves what happened, why it happened, and how to sustain strengths and improve on weaknesses" (p. 1). In other words, the units perform a collective self-examination in which the more general question, "How did the unit do?," is broken down into three more specific questions:

1. "What happened during the collective training exercise?" In other words, AAR participants attempt to specify the facts (i.e., the important actions and outcomes) of the simulated battle.

2. "Why did it happen?" Given the facts of the exercise, the participants attempt to explain the causes of particularly important actions and outcomes.

3. "How can units improve their performance?" Given that the previous two questions are answered, the participants determine appropriate actions to solve problems identified in their performance."

While these reviewers assert that this procedure "enables soldiers to discover for themselves what happened," the AAR as typically performed is lead, managed, and directed by an external agent, not the individual soldiers.

Morrison and Meliza go on to link key elements of the AAR to experiential learning. "This approach stipulates that learning is facilitated by real-world experiences and is therefore often associated with the phrase 'learn by doing.'"

And, "Perhaps the most basic principle on which the AAR is based is that learning and performance are enhanced when appropriate feedback is provided. Feedback refers to information that people receive during or after performance of an action to control and learn the action."

That is to say that feedback is provided by an "other."

Morrison and Meliza do go on to concede that there is reason to suspect that the AAR is not always performed with the objectivity and care that might be desired. This likelihood is a primary basis for the alternative explored in this paper.

AAR HISTORY

I/ITSEC has a history of interest in the AAR. TNO in the Netherlands (Buiëll and Lubbers, 2007) have discussed companion agents that would support reflection. There has been considerable interest in the United States in adding automated features as well, notably by Stottler-Henke (Chen, Jensen, Bascara and Harmon, 2007), (Stottler, 2003), (Jensen, Chen, Nolan, and Jacobs, 2005), (Chen, Jensen, Bascara, and Harmon, 2007), and BBN (Travers, Ferguson and

Langevin, 2007), as well as others (Ekblad, Gonzalez, Fernlund, and Barath, 2005), (Frank, Whiteford, Hubal, Sonker, Perkins, Arnold, Presley, Jones and Meeds, 2004), and (Kelly, 1999).

Across all of the manifestations and manipulations of the AAR, however, a singular flaw can be detected. The conclusions of the AAR inevitably depend to a greater or lesser extent on the analytical perspective of someone whose judgment always can be challenged. The "official" result comes from an official, not from those who participated in the experience. This can result in an alienation of ownership of outcomes – as Piaget noted in research with game-playing children in the Swiss Alps almost 100 years ago.

Piaget observed that "Most of the rules the child learns to respect [in typical discourse with the world] he receives from adults, which means that he receives them after they have been fully elaborated, not in relation to him as they are needed, but once and for all and through an uninterrupted succession of earlier adult generations.

"In the case of the very simplest social games, on the contrary, we are in the presence of rules which have been elaborated by the children alone."

Further, "In short, law now emanates from the sovereign people and no longer from the tradition laid down by the elders." (Piaget, 1965.)

The children 'own" their experience during gameplay in ways not always associated with traditional forms of education.

At its best, the AAR would boast the benefits of a good narrative in the making of meaning, and it would do so in a way that also makes for an owning of the meaning.

Jerome Bruner is, quite possibly, the individual most known for articulating the ways that narrative processing is nothing less than the active process of meaning-making (Bruner, 1990, 2003). What is needed with an ideal AAR, then, is a capability that allows for the making of meaning without breaking cognitive engagement with the experience.

MODEL-BASED APPROACHES TO NARRATIVE IN AFTER-ACTION REVIEW

For the remained of the paper, we consider the construction of an AAR system that approaches that goal. This is a system built within a 3D virtual environment in which units perform actions to achieve mission goals, interacting with other human-controlled units, as well as system resources that model the environment and other agents. When considering how such a system might be developed to provide effective AAR, it is critical to keep as design goals the three components of AAR referred to above, namely the effective explication of *what* happened, *why* it happened and *how* the involved units can improve their performance in subsequent challenges.

Content determination: What happened

A primary task for an automated AAR system is to determine the content of the review – just what portion or parts of the users' actions does the system pull out during the review for discussion? Systems that log all user activity can begin to answer this question by using the log to rewind or fast-forward through a user's history. These kinds of systems require very little in terms of action representations to perform effective logging. Typically, code that runs each action within a simulation is augmented with a database update, tagging an action execution with a timestamp and other relevant information. Use of logs within an AAR session, however, require a human reviewer to drive the review and are subject to the limitations of human-driven AAR as described earlier.

In contrast to a human-driven log review, a *model-based* approach, with an explicit model of activity within the simulation domain, can provide a level of representation that allows automated reasoning algorithms drawn from artificial intelligence research to select review content. One such model is that used by the Zocalo system (Young, et al, 2002), designed for the intelligent control of virtual worlds like computer games and training simulations. The Zocalo system is currently being used by the Liquid Narrative Group at North Carolina State University to construct interactive narrative-oriented games. The architecture is specifically designed to bridge the gap between game and simulation engine design/development and work in artificial intelligence that focuses on the automatic creation of novel and effective action sequences.

The process of constructing action sequences – or plans – for execution within a virtual environment involves a number of specialized functions, including reasoning about the actions of individual characters, generating any character dialog or narration to be provided by the system, and creating cinematic camera control directives to convey the action that will unfold

in the story. To facilitate the integration of corresponding special-purpose reasoning components, the Zocalo architecture is highly modular. Individual components within Zocalo run as distinct processes (typically on distinct processors, though this is not a requirement); components communicate with one another via a well-defined XML-based message-passing protocol; developers extending Zocalo to provide new functionality wrap their code within a message-passing shell that requires only a minimal amount of customization.

Adopting a plan-based model of story structure allows the system to compose new stories in response to novel starting states or goal specifications, or to customize a story based on a user's interests and knowledge. The use of a formal plan representation has two additional advantages. First, the formal properties of the planning algorithm guarantee that the plans contain adequate structure to effectively control the story world's virtual environment. Specifically, the plans that Zocalo generates are provably sound, that is, when executed, each action in them is guaranteed to execute correctly and the plans themselves are guaranteed to achieve their top-level goals. These properties make the plans it produces well-suited for use in controlling the execution of a virtual environment.

A second benefit to the use of plans to drive a narrative is in the plan's structural correspondence to a user's mental model of the story it defines. Our recent research (Christian and Young, 2005) suggests that hierarchical causal link plans like those used within Zocalo, as well as the techniques used by Zocalo's planning algorithm to create them, make for effective models of human plan reasoning. Our empirical studies indicate that the core elements of our plans match up with the models of narrative structure defined and validated by psychologists (Lauer, Peacock, and Graesser, 1992). By using a formal representation for story structure that corresponds to users' models of stories, we can make more direct predictions about the users' understanding of the stories we create. We rely on this correspondence when designing techniques to create specific narrative effects.

Using Zocalo, system developers construct two parallel models of the game world, one using extensions to existing game engine code, the other using techniques for explicit modeling of actions in terms of their requirements for execution and their effects on the game world. When integrated with a virtual world engine, Zocalo acts as a run-time behavior generator, responsible for both generating plans – coherent action sequences that achieve a specific set of in-game goal – and maintaining the coherence of those plans as they execute in the face of unanticipated user activity.

In Zocalo, simulator actions have a dual representation. One the one hand, Zocalo uses a *procedural* representation of an action, containing the code needed to make an action execute within a virtual environment. On the other hand, Zocalo employs a *declarative* representation of an action, explicitly stating the context and consequences of an action's execution. The declarative representation builds on a plan-based approach first developed by Fikes and Nilsson (1977) in which each action is defined by an explicit set of *preconditions* – a list of all the conditions in the simulation world that must be true in order for an action to execute correctly – and a set of *effects* – the complete list of ways that the action's execution will change the environment once it executes successfully.

Within Zocalo, system designers build the actions available to their users and system-controlled characters using a tool called Bowman (Thomas and Young, 2006) that provides support for the specification of declarative definitions of *action libraries.* Once a collection of actions has been defined in Bowman, the tool automatically translates the action definitions into procedures, that is, software capable of implementing each action's execution within a target virtual world. Zocalo's dual-representation system is designed to facilitate a broad class of reasoning about action within the virtual world (some of which is described below) that can directly be translated into code that can control the environments it models. The relevant benefit to its use is that logs of actions executed in a Zocalo environment allow the automatic piecing together of complicated causal chains between user- and system-performed actions that would be hard to re-construct in a system that relied on a procedural representation (and the logs it can generate) alone.

While AARs are not always viewed as stories, the process of constructing an AAR system's report of *what* happened often shares much with the process of constructing the action-oriented contents of a narrative. Both narratives and AARs contain descriptions of action sequences and are selected by an author in order to convey information about the underlying meaning within the actions themselves. At a basic level, this selection process centers around causality and temporality: what happened, when did it happen, and how did it change the world in ways that are important for later actions. The first and second of these questions may be answerable by reviewing execution

logs alone, but a model-based approach is needed in order to provide answers for the third question.

To answer questions about the causal significance of actions, a system like Zocalo can infer or reconstruct the causal relationships between actions based on the declarative representation's explicit marking of the preconditions and the effects for every action in an execution trace. By matching an early action's effect with the precondition of a later action, one can readily identify the later action's dependency on the earlier act's successful execution. By identifying actions with effects that undo conditions established by earlier actions, one can identify flaws in an execution sequence that explain why some actions fail to execute correctly.

The process of taking a Zocalo execution log and reconstructing the complete causal network covering all of its actions mirrors that of plan generation used by AI planning systems. Plan generation involves the construction of action sequences from representations of individual actions such that the sequences are guaranteed to move an agent from a known current state to a given goal state. In recent work, we have combined this plan re-construction capability with results from work in Cognitive Psychology on the role of causality in narrative comprehension to build effective game summaries from Zocalo game logs (Cheong et al, to appear; Cheong and Young, 2006) .

Causality alone, however, cannot provide the complete answer to the "why it happened" question. For instance, causal relationships between actions do not describe a user or character's *choice* between alternative actions, a user's focus of attention when observing the simulation world, or his or her missing domain knowledge.

Content determination: Why did it happen?

Determining how to explain why an action or sequence of actions has happened is a complicated process even for a human expert. To address this problem, we again refer to the Zocalo action representation, and methods we have previously developed for the generation of narrative-based action sequences in games. The *why* behind an action can be explained in part at the causal level but also at the intentional level. We have described above how an AAR system using Zocalo might re-construct the causal network characterizing a training session's actions, and this network can provide some information about the process. However, this representation fails to take into account the motivations and intentions behind character and/or player actions, an important aspect of explaining why events transpire in a particular manner.

How can units improve their performance?

One often-used way of explaining how units can improve their performance in AAR is by explicit instruction. In this situation, an instructor opportunistically refers to points during an exercise where units performed problematically. While this approach proves useful when scenarios expose appropriate problems, when units perform well, this type of AAR has essentially no grist for the mill. Instructors in AAR are at the whim of the units and their behavior when referring to problematic points in exercises.

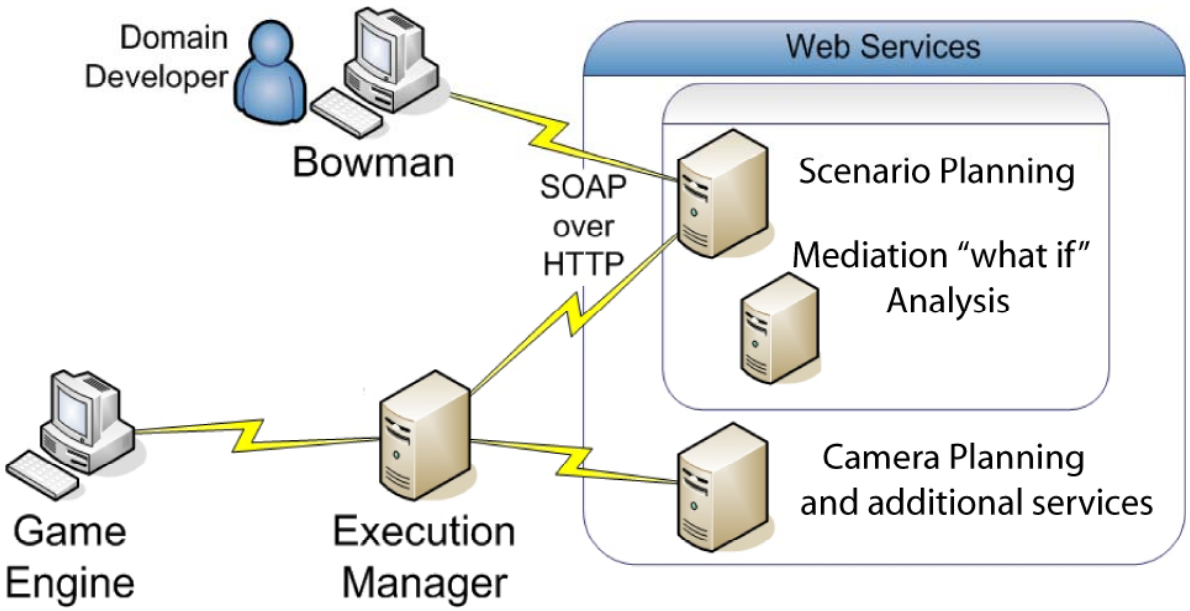

Figure 1. The Zocalo Architecture.

A potential alternative method that Zocalo's model-based approach allows is one where scenarios are generated by the system with the intention of exposing misconceptions and/or facilitating problematic behavior, in order to provide specific learning opportunities. The typical model for Zocalo's current use is to generate plot lines for a game level that are composed of plans for the action of both a user and all system-controlled characters. These plot-lines are constructed by referring to the action library and actions are added based on the system's need to establish all preconditions of actions prior to their point of execution.

By providing the plot-line construction process with an explicit *user model* characterizing potential misconceptions, the same plan construction methods can be used to search for plans that contain potential *flaws*. These flawed plans would correspond to courses of action in which units would, due to a lack of knowledge about some element of the task domain, select incorrect courses of action and consequently perform poorly in the given scenario.

For instance, knowing of a particular knowledge deficit of a unit commander, the system could design a challenge requiring the commander to make a choice for action based on his or her incorrect understanding of the domain. This would result in some form of action failure. In the face of this failure, and armed with the model-based understanding of the reasons behind it, an AAR system could provide contextualized feedback in the moment. This sort of instructional intervention prompts metacognitive reflection-on-action during the course of the action, rather than halting it. Such *self-learning* strategies have been shown to be very effective (Chi, 2008) in establishing continuously situated meaning. (Of course, the system also will maintain the record as a summary for use during the conventional discussion after the completion of the action if that is desired.)

EXPLORING A MODEL-BASED APPROACH FOR PROVIDING FEEDBACK IN AAR

We have provided a general sketch above for ways that model-based approaches to the control of action within a virtual world may be used to structure training scenarios. The declarative models of actions and plans that lie under the scenarios provide automated reasoning tools the knowledge base needed to identify problematic behavior. Further, these models can allow a scenario generation system to explore the space of scenarios to select ones that are most helpful at providing challenges to a user's current skill levels.

In recent research, we have used these same models to move beyond the identification of actions and plans to the automatic description of actions via text and cinematic 3D video. For example, using the same plan-based models of action that are used in Zocalo, we have generated instructional text from the data structures describing complex tasks (Young 1999a).

The efficacy of these instructions has been tested by verifying that users following these instructions within a virtual environment perform their tasks with fewer errors and achieve more of their goals than users following instructions generated by several competing methods (Young 1999b).

In more recent work, we have built a layer on top of our action models that represents the knowledge used by cinematographers when controlling a camera filming complex action sequences (Jhala and Young, 2005). Combining this model with previous work on the generation of natural language discourse (Young, Moore and Pollack, 1994), we have built several systems that take as input descriptions of action sequences and produce cinematic video describing the action execution within a game environment (e.g., Jhala, Rawls and Young, 2008).

Because both the text and cinematic generation systems build upon Zocalo's underlying action model, it is possible to extend this work so that both text and video providing explanations for action sequences and their failures could be automatically created based on execution logs used for AAR.

CONCLUSION

Under ideal conditions, an After Action Report is often key to realizing an understanding of *what* happened and *why* it happened and *how* new behaviors and knowledge can be learned for future benefit. Yet this system, with fine-grained learner tracking that can precipitate teaching moments – and through "graceful failures" preclude disruption of the experience, offers new possibilities for maintaining active cognitive engagement in learning.

ACKNOWLEDGEMENTS

Support for Young's work described in this paper has been provided by National Science Foundation CAREER award # 0092586.

REFERENCES

Bruner, Jerome. (1990). *Acts of Meaning.* Boston, MA:Harvard University Press.

Bruner, Jerome. (2003). *Making Stories: Law, Literature, Life*. Boston, MA:Harvard University Press.

Buiëll, E. F. T., and Lubbers, J., (2007). A companion agent for automated training systems, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Chen, D. Y., Jensen, R. Basscara, O., and Harmon, N. (2007). Enabling automated AAR development by abstracting data collection from analysis, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Cheong, Y. and Young, R. M. (2006). A framework for summarizing game experiences as narratives, Proceedings from the Second Conference on Artificial Intelligence and Interactive Digital Entertainment, Stanford, CA:Stanford University.

Cheong, Y, Jhala, A., Bae, B. and Young, R. M. (to appear in 2008). Automatically generating summary visualizations from game logs, Proceedings from the Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 08), Stanford, CA: Stanford University.

Chi, M. (2008). Self-learning. Invited Speech. Advanced Distributed Learning Speaker Series, Alexandria, VA:Advanced Distributed Learning Initiative.

Christian, D. and Young, R. M. (2004). Comparing cognitive and computational models of narrative structure. Proceedings from the National Conference of the American Association for Artificial Intelligence. Menlo Park, CA:AAAI.

Edblad, J., Gonzales, A., Fernlund, H. and Barath, P. (2007). Automatic detection of discrepancies in after action review, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Fikes, R. E. and Nilsson, N. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, 5(2):189-208.

Frank, G. Whiteford, B., Hubal, R., Sonker, P., Perkins, K., Arnold, P. Presley, T., Jones, R., and Meeds, H. (2004). Performance assessment for distributed learning using after action reports generated by simulations, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Jhala, A., Rawls, C. and Young, R. M. (2008). Longboard: A sketch-based intelligent storyboarding tool for creating machinima, Proceedings from the Florida Artificial Intelligence Research Society Conference. Orlando, FL:Florida Artificial Intelligence Research Socieity.

Jhala, A. and Young, R. M. (2005). A discourse planning approach for cinematic camera control for narratives in virtual environments. Proceedings from the National Conference of the American Association for Artificial Intelligence. Menlo Park,

CA:Association for the Advancement of Artificial Intelligence.

Kelly, M. (1999). After action review in synthetic environment based training systems: A training system not a technology, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Kennedy, M. M., (1987). Inexact Sciences: Professional education and the development of expertise, Issue Paper 87-2, *Review of Research in Education*, XIV. Washington, D.C.:American Educational Research Association.

Knerr, B. W., Lampton, D. R., Martin, G. A., Washburn, D. A., and Cope, D. (2004). Developing an after action review system for virtual dismounted infantry simulations, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Lauer, T., Peacock, E., and Graesser, A. C. (1992). (eds.) *Questions and Information Systems*. Hillsdale, NJ:Erlbaum

Mastaglio, T. W., Jones, P., Bliss, J. P., and Newlon, E. (2007). An integrated theory for after action review, Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Piaget, J. (1965). The rules of the game of marbles, in Bruner, J. S., et al. (eds.) (1976). *Play: It's Role in Development and Evolution.* New York, NY:Basic Books.

Riedl, M., Saretto, C.J. and Young, R. M. (2003). Managing interaction between users and agents in a multiagent storytelling environment, Proceedings from the Second International Conference on Autonomous Agents and Multi-Agent Systems. A Melbourne, AU:International Foundation for Autonomous Agents and Multi-Agent Systems.

Riedl, M., and Young, R. M. (2005). An objective character believability evaluation procedure for multi-agent story generation systems. Proceedings from the International Conference on Virtual Agents. Kos, GR: International Society for Virtual Agents.

Squire, K., (2006). From content to context: Videogames as designed experience. *Educational Researcher*, Vol. 35., No. 8. Washington, D.C.: American Educational Research Association.

Stottler, R. (2003) Techniques for automatic AAR for tactical simulation training. Proceedings from Interservice/Industry Training, Simulation and Education Conference. Arlington, VA:NTSA.

Thomas, J. and Young, R. M. (2006). Elicitation and application of narrative constraints through mixed-initiative planning. Proceedings from the International Conference on AI Planning Systems Workshop on Preferences and Soft Constraints in Planning. Menlo Park, CA:International Association for the Advancement of Artificial Intelligence.

Young, R. M., Mark O. Riedl, Mark Branly, R.J. Martin, C.J. Saretto. (2004). An architecture for integrating plan-based behavior generation with interactive game environments, *Journal of Game Development*, Vol. 1., No. 1. Boston, MA:Charles River Media.

Young, R. M., Moore, J. D., and Pollack, M. E. (1994). Towards a principled representation of discourse plans, Proceedings from the Sixteenth Conference of the Cognitive Science Society. Boston, MA:Cognitive Science Society.

Young, R. M. (1999a). Using Grice's maxim of quantity to select the content of plan descriptions, in *Artificial Intelligence*, Vol. 115, No. 1. Boston, MA:Elsevier Press.1999a.

Young, R. M. (1999b). Cooperative plan identification: Constructing concise and effective plan descriptions. Proceedings from the National Conference of the American Association for Artificial Intelligence. Menlo Park, CA:Association for the Advancement of Artificial IntelligenceOrlando, FL. August.