Measuring Presence and Performance in a Virtual Reality Police Use of Force Training Simulation Prototype

Edward T. Garcia, Stephen G. Ware
University of New Orleans
2000 Lakeshore Drive
New Orleans, LA, 70148, USA
etgarcia@uno.edu, sgware@uno.edu

Lewis J. Baker
pymetrics Inc.
102 Madison Ave
New York, NY, 10016, USA
lewisjbaker@gmail.com

Abstract

In this paper we describe a virtual reality training simulation designed to help police officers learn use of force policies. Our goal is to test a training simulation prototype by measuring improvements to presence and performance. If successful, this can lead to creating a full-scale virtual reality narrative training simulation. The simulation uses a planner-based experience manager to determine the actions of agents other than the participant. Participants' actions were logged, physiological data was recorded, and the participants filled out questionnaires. Player knowledge attributes were authored to measure participants' understanding of teaching materials. We demonstrate that when participants interact with the simulation using virtual reality they experience greater presence than when using traditional screen and keyboard controls. We also demonstrate that participants' performance improves over repeated sessions.

Introduction

Simulations provide a safe environment where trainers can teach trainees about situations that may be dangerous. Simulations can be more cost-effective than using actual equipment, both in terms of training costs equipment repair. Military personnel, firefighters, doctors, nurses, and police employ simulations for training. (Hays and Singer 1989)

We have created a prototype simulation that allows police officers to explore the consequences of use of force decisions. It takes place in a virtual world where a police officer responds to a call about a potentially dangerous suspect. The participant takes the role of the officer and is free to take any action available in the virtual world. The goal is to teach key concepts defined by the Police Executive Research Forum 2012 via a learn-by-doing approach. When relying on human actors, these kind of role-playing scenarios can be time-consuming and cost-prohibitive to create, run, and evaluate. Even our limited scope prototype domain, which contains 11 types of actions, 5 endings, and 5 measures of player knowledge, can be in 125,688 unique states and allows 752,741 possible transitions between those states. Hence we use a planner-based experience manager in place of human actors to make decisions in our simulation.

In our evaluation of this system, we logged event data, tracked indicators of player knowledge, measured physiological signals, collected participant self-reports about their experience, and used two methods of scoring performance to understand the factors that influence presence and learning in the simulation. Our contribution is the development of an intelligent tutoring system that uses an automated experience manager to create interactive stories that teach use of force policies, and we show that using virtual reality can improve presence. We expect this increase in presence to lead to increased transfer of knowledge (Alexander et al. 2005).

Related Work

Training and Presence

Various studies have suggested ways to measure training effectiveness. We are ultimately interested in *transfer*, the application of knowledge in a real world environment that was learned in a virtual one. Several factors may affect transfer (Alexander et al. 2005). This study focuses specifically on *presence* (Slater and Steed 2000), the experience of being in one place when one is actually in another. We want to test whether newly available, affordable, commercial virtual reality technology can increase presence. Stevens and Kincaid (2015) point out a clear link in the literature between virtual reality training and transfer. Though the relationship between presence and transfer is less clear, they demonstrated a moderate relationship between presence and performance in a virtual environment, with the expectation that improving presence can improve transfer. Because our prototype simulation is not yet in a position to demonstrate transfer (for example, through fewer incidents of unjustified use of force in a community), we will measure presence and performance as a proxy for effectiveness.

Specifically, we use Witmer and Singer's (1998) Presence Questionnaire, which has demonstrated high inter-rater reliability and has been widely used. It measures 4 factors relating to presence: Control, Sensory Experience, Distraction, and Realism. We also use physiological measures suggested by Riva, Davide, and IJsselsteijn (2003), who demonstrated a link between presence and physiological signals like heart rate, skin conductivity, and skin temperature.
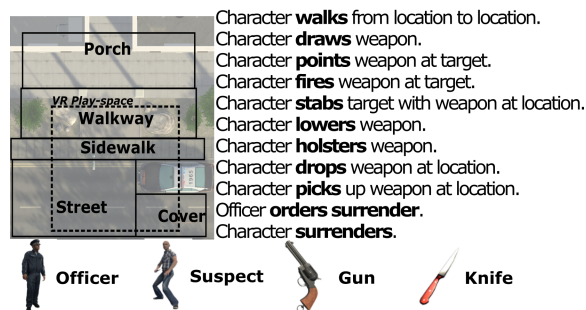
Figure 1: This map shows how physical space in our domain was discretized and the characters, item, and actions available in the simulation.

Narrative Planning

Planning is a branch of AI research that reasons about a sequence of actions which achieves a goal. A classical planner (Newell, Shaw, and Simon 1959) takes as input (1) a description of the world in some formal logic, (2) a group of action templates with preconditions and effects, and (3) a goal. It generates a sequence of executable actions (i.e. a plan) to achieve the goal.

Narrative planning (Young et al. 2013) expands upon classical planning by adding additional constraints to ensure that agents act in a believable way. Some system-wide goal called the author's goal must be satisfied. Agents also have individual goals, and every action they take must be in service of one of those goals. Agents may conflict or cooperate as they pursue their goals, while the planner as a whole shepherds all agents toward the author's goal (Ware 2014). Riedl and Bulitko (2013) classify narrative planners as *strong story* systems, which may be preferable to *strong autonomy* or *simulation-based* training systems when the designer places strong constraints on the content and outcome of the narrative, as we do in this system.

Experience Management

An interactive narrative must be prepared for the user to take any action that is offered to them. Typically, each user controls a character (their avatar) while an experience manager controls all non-player characters and the environment. Roberts and Isbell (2013) provide a survey of experience management techniques. The experience manager's goal is to increase the quality of a participant's experience as defined by a human author who specifies certain criteria, principles, goals, and aesthetics (Riedl et al. 2008). For our simulation, we want users to learn about use of force decisions by experiencing positive endings when they apply policies correctly and negative endings otherwise, all in a safe and low-risk environment.

Intelligent Tutoring Systems

Intelligent Tutoring Systems (ITS) can employ narrative planning and experience management to create environments that engage and teach at the same time. VanLehn (2006) and Ma et al. (2014) provide surveys of ITS research.

Specifically relevant to this project are systems like *Annie* (Thomas and Young 2010) and *Automated Scenario Adaptation* (Niehaus, Li, and Riedl 2011), planning frameworks for training that automatically generate and proactively adapt scenarios for users based on a plan-based description of the task environment. Many ITS also use the learn-by-doing approach we have adopted (e.g. Schank, Berman, and Macpherson 1999, Zook et al. 2012, and many others). Our prototype is distinct from many previous plan-based systems because it considers the long-term consequences of each decision on the space of possible stories, allowing it to balance virtual character believability and pedagogical structure while allowing the player extensive freedom.

Simulation Description

The simulation domain consists of a police officer (Officer), who is responding to a call from a young man's mother. The young man (Suspect) has been recently kicked out of the house, is on the porch, banging on the door, and possibly has a knife. The participant plays the role of the police officer.

Architecture

The domain has been written in STRIPS-style format where all objects, actions, and goals are specified (Fikes and Nilsson 1972). The simulation uses a client/server configuration, where the client takes input from the participant to control the officer, and the server controls all other agents. The client was created with the Unity Game Engine. The play-space and actions taken by the client are discretized for processing by the planner on the server. The server stores the current state of the simulation, and depending on the state, the participant is allowed to take certain actions. The server also directs the client on which actions the Non-Player Characters (NPCs) will take. The client can be run using either Screen and Keyboard controls or Virtual Reality controls.

The virtual reality controls use the HTC Vive virtual reality system in room-scale mode, where the room represents the area shown in Figure 1. The Vive is a motion-tracking headset which displays a visual representation of the virtual world and two motion-tracking hand controllers that represent the participant's hands.

Our simulation contains 5 locations, 2 characters, 2 weapons, 11 actions, and 5 axioms, shown in Figure 1. Each session begins with the Officer at the Cover location and the Suspect at the Porch location. A session is designed to last about 1 minute. Each action in the simulation takes roughly the same amount of time as it would take in the real world. Every successfully executed character action (by the player or an NPC) is logged for analysis. At the end of each session a score is displayed on the screen depending on how the simulation ended. Scores rank the possible endings from worst to best based on the safety of the officer and the suspect.

- Score 0: The suspect stabbed and killed the officer.
- Score 1: The officer shot the suspect, but the suspect never threatened the officer with the knife.
- Score 2: The officer shot the suspect after being threatened with the knife.

- Score 3: The suspect surrendered, but only after threatening the officer with the knife.
- Score 4: The suspect surrendered and never threatened the officer with the knife.

Use of Force Policies

The Police Executive Research Forum (PERF) (2012) has identified best practices regarding use of force that are designed to ensure that the officer, suspect, and bystanders remain safe in a potentially dangerous situation. PERF also speculates that some officers may leave the academy with a bias toward using force because many training simulations assume that force is always necessary. There are many tools for teaching officers how to shoot, but too few for teaching them how *not* to shoot. They call for innovative methods to address this problem.

Our simulation is designed to teach one use of force policy in particular: *distance + cover = time*. When an officer keeps distance and cover between himself and a suspect, he can buy time to achieve a peaceful resolution. Policies like this one demonstrate the advantages of interactive narrative training simulations over traditional shooting range simulations because, depending on the trainee's actions, force may not be needed at all.

We determined five specific features about which trainees might demonstrate knowledge or ignorance. We call these *player knowledge attributes*, and define them based on the states in which the trainee finds himself and the actions he takes or does not take. Knowledge or ignorance of these attributes can thus be measured automatically by analyzing a session log. We must note that these represent our own non-expert interpretations of use of force policies. Before using this simulation to train actual police officers we must obtain feedback from experts.

- Keep Distance: The officer should keep distance between himself and the suspect. If the participant and suspect get within arm's reach of one another (that is, occupy the same location as shown in Figure 1), this concept is not known. Otherwise, this concept is known.
- Use Cover: The officer should keep cover between himself and the suspect, even if he must retreat. If the participant walks back to the Cover location (behind the car), this concept is known. Otherwise, this concept is not known.
- Justified Force: PERF (2012) mentions that holding a knife is not the same as brandishing a knife. If the suspect raises the knife in a threatening way and the officer uses deadly force, this force was justified. If the officer finds himself in a situation where force is justified and uses it, this concept is known. If the officer finds himself in a situation where force is justified and does not use it, this concept is not known.
- Unjustified Force: If the officer used deadly force when the suspect was not close and/or has not raised the knife, force was not justified and should not have been used. If the officer uses force in this way, this concept is not known. Otherwise, it is known.

- Agitation: The suspect is nervous, and the way the officer deals with him can either calm him down or further agitate him. If the officer points his gun at the suspect while the suspect is not angry, the suspect becomes angry and aggressive. If the officer angers the suspect in this way, this concept is not known. Otherwise, it is known.

Narrative Control

In the simulation, the player should always be free to perform any action when its preconditions are met. By using an automated experience manager, our system is able to react to a player taking any of these actions without needing to hand author every possibility. In some cases, the experience manager may direct an NPC to take an action. In other cases, the experience manager may decide it best for the NPC not to act at all.

The state space of the simulation can be represented by a story graph where nodes are unique states and a directed edge $s_1 \xrightarrow{a} s_2$ may exist from state s_1 to state s_2 for action a if a is allowed in state s_1 and taking a in s_1 would results in s_2. Each action can be categorized as either a player action or an NPC action.

A state is a set of propositions which completely describes three things: the configuration of the physical world, the intentions of all agents, and the current state of the five player knowledge attributes. Knowledge attributes can be known, not known, or unobserved.

The experience manager's decisions have been precompiled using the methods below, but the same criteria could be applied (or approximated) in realtime systems. We begin with the full story graph for our domain—the entire state space representing every state and every possible action reachable from the initial state. We then prune this graph intelligently until every NPC has at most one action to perform in each state, thus making the experience manager's decisions unambiguous. We never prune player actions (i.e. we never prevent a player from taking an action which should be possible in the current state). Pruning the story graph at design time allows us to fully consider the long-term consequences of every decision on the space of possible stories that can be told. The full story graph for this domain contains 125,688 nodes and 752,741 edges.

Intentional Prune Studies show that virtual characters appear more believable when they act intentionally—that is, they appear to be working toward their goals (Riedl and Young 2010). We use Riedl and Young's model of intention: Any action that cannot contribute to achieving an agent's goal should be pruned. After intention pruning, the story graph is reduced by 0.2% nodes (125,428) and 26.4% edges (554,319) from the full story graph.

Unique Ending Prune The author's goal in this domain is a disjunction of various possible ending states. The experience manager is neither cooperating with the player to achieve a good ending nor opposing the player to achieve a bad one. Rather, the actions taken by the player (not the NPCs) should be responsible for the ending earned. Hence NPCs should prefer actions which keep the higher number

of possible endings available. This is a tie breaking prune, which means that if there exists only one edge for an NPC, it will not be pruned using this technique. After unique ending pruning, the story graph is reduced by 0% nodes (125,428) and 0.7% edges (550,447) from the previous graph.

The order in which these pruning strategies are employed is important. We perform intention pruning first because it is important for characters to be believable. If unique ending pruning were to happen first, it is possible that the characters would act unbelievably or not at all to ensure more endings stay available. For example, say the Officer angers the Suspect. The Suspect approaches the Officer and then raises the knife. If unique ending pruning has occurred before intentional pruning, the Suspect would simply do nothing from this point on, because stabbing the officer would remove a unique ending. Instead, we want the Suspect to follow through with his plan, even if it reduces which endings are available.

Player Knowledge Prune   When the simulation starts, each player knowledge attribute is set to unobserved. These attributes are the simulation's model of player knowledge. Given multiple NPC actions, the action that leads to observing a player knowledge attribute (whether known or unknown) as quickly as possible is preferred. This is also a tie breaking prune; if there exists only one NPC edge, it will not be pruned. After player knowledge pruning, the story graph is reduced by 0% nodes (125,428) and 1.1% edges (544,491) from the previous graph.

We prioritize keeping unique endings available over learning about the player. If we had done the reverse, the following example could have occurred: The Officer angers the suspect, so Suspect immediately raises the knife. This is the quickest way for the simulation to determine if the officer will use justified force. However, this also immediately eliminates all endings where the Suspect did not threaten the Officer, making it impossible to achieve the best score. Instead, the Suspect approaches the Officer first, then raises the knife, keeping more unique endings available and still allowing the simulation to learn if the Officer will use justified force.

Arbitrary Prune   At this point, there may still be a few states that have multiple possible NPC actions. We treat all of these actions as equally good. Consequently, we arbitrarily prune by choosing the first action. After arbitrarily pruning, the story graph is reduced by 0% nodes (125,428) and 1.7% edges (534,991) from the previous graph.

Evaluation

We tested whether virtual reality technology increases a trainee's experience of presence and whether this simulation teaches use of force policies. Twenty-one civilians were recruited to test the simulation. We tested with civilians because we did not want to risk adverse effects if the simulation failed and because the concepts being taught did not require specialized knowledge or police training.

Methodology

We scheduled participants to come into our virtual reality lab. Before starting the simulation, each participant was as-

signed a number for anonymity and watched a video about safety and the controls for the simulation. Due to the novelty of virtual reality technology, subjects completed a short tutorial that required them to use all of the simulation's controls and were also given a verbal quiz about the controls. No detail was given about the content of the simulation, what we were measuring, or what behavior we expected from the subjects. Investigators were prohibited from answering subjects' questions about what they should and should not do; they were only permitted to answers questions about safety and the simulation's controls.

We alternated the starting control type so that half of the participants started with screen/keyboard controls, while the others started with virtual reality controls. At the end of each session, the participant was shown a score between 0 to 4 inclusively (as defined earlier). The participant was required to play at least 2 sessions and was allowed to play up to 10 sessions per control type. The participant then played the simulation using the other set of controls, and was again allowed to play between 2 and 10 sessions.

Virtual Reality Improves Presence

We collected self-report questionnaires at several stages of training as well as participants' physiological state monitored via wrist-strap sensor throughout all sessions of the simulation. Physiological and self-report measures of presence are highly correlated (Meehan 2001), and both give some insight into participants' experience of virtual reality versus screen and keyboard controls.

Self-Report Measures of Presence   Before the simulation began, participants completed Witmer & Singer's immersive tendencies questionnaire (ITQ) to measure baseline experience of presence. After running the simulation using the first control type, the presence questionnaire (PQ) was given to measure presence in virtual environments (Witmer and Singer 1998). Subjects completed a modified presence questionnaire (MPQ) at the end of the study, where they chose whether they preferred virtual reality, screen/keyboard, or no preference for each factor in the PQ.

Participants were randomly assigned to begin with either virtual reality or screen/keyboard controls (N = 21, 10 VR first and 11 screen/keyboard first treatments). A Welch's independent t-test revealed that participants who played VR first reported marginally higher overall presence in the PQ ($t_{16.32}$ = -1.664, $p = 0.1152$).

Participants completed the simulation using both virtual reality and screen/keyboard controls. Looking at self-reports of preference in the MPQ, participants strongly preferred virtual reality controls (x_2^2 = 12.373, $p = .002$, $\varphi = .142$). Additionally, participants who initially used screen and keyboard were 10.8% (S&K first: 73% and VR first: 62.3% prefer VR) more likely to prefer virtual reality across all items of the MPQ ($x_1^2 = 7.396$, $p = .007$, $\varphi = .136$).

Participants completed the PQ after their first experience with the simulation (using either set of controls). The PQ therefore measured presence in the screen and keyboard version for 11 participants, and virtual reality for the other 10 participants, limiting comparison between groups. A logis-

Figure 2: Average heartrate, skin conductivity, and skin temperature of participants progressing through the simulation.


Figure 3: Left: Average score of all participants as they progressed through the simulation. Right: Overall average knowledge attribute scores for all subjects as they progressed through the simulation.

tic regression analysis revealed no relationship between reported presence on the PQ by control preference indicated in the MPQ (z = -0.516, p = .606).

It was possible that individuals with high baseline presence might differ in their experience of presence and control preference in the simulation. An ANOVA modeling PQ scores by ITQ and treatment revealed marginally significant PQ scores for participants who played VR first ($F_{1,18}$ = 4.079, p = 0.060, η^2 = .165), and marginally significant PQ as ITQ increased ($F_{1,18}$ = 3.354, p = 0.085, η^2 = .169), with no interaction ($F_{1,18}$ = .081, p = 0.381). There was no difference in overall preference for VR by ITQ scores ($F_{1,20}$ = 1.442, p = 0.245).

Physiological Measures of Presence An Empatica E4 physiological sensor was attached to the participant's right wrist to measure heart rate, skin conductivity, and skin temperature throughout every simulation session. Meehan (2001) proposed the following relationships between these measures and presence. Under stress, heart rate increases, skin conductivity increases, and skin temperature at core increases which means decreased skin temperature at the extremities. We expect the same results as the participant experiences higher presence.

Figure 2 visualizes a mixed-effect multilevel model of physiological response within participants across multiple sessions and controls, and accounting for whether participants first played VR or keyboard controls. This analysis revealed that controls significantly increased heart rate (t_{131} = 4.244, p < .0001) and skin conductivity (t_{132} = 4.111, p < .0001), and significantly decreased skin temperature (t_{131} = -2.70, p = .008) when playing the virtual reality simulation, all indicative of increased presence.

Training Effectiveness

We measured performance in two ways. *Score* is a single value that represents performance in a particular session. It is meant to rank the possible endings, with higher scores representing better endings. Score was shown to subjects after each session. *Player knowledge attributes* represent specific training concepts of which subjects could demonstrate knowledge or ignorance. A subject's overall knowledge is represented as a vector of nullable boolean values. These attributes were never described to subjects, and their performance on them was not shown to them.

Score A mixed-effect multilevel model predicting scores from treatment group, controls used, and session number

nested within participant revealed that participant scores increased over repeated sessions (t_{145} = 4.609, p < .0001) as shown in Figure 3.

Player Knowledge Attributes As an alternative measure of performance, we automatically analyzed each subject's session log based on the five player knowledge attributes identified earlier. Each component of the player knowledge vector can be represented with the following values:

- Unobserved: Simulation does not have enough information to determine whether the knowledge attribute is known.
- Known: Simulation has observed an action that indicates the subject knows the attribute.
- Not Known: Simulation has observed an action that indicates the subject does not know the attribute.

A subject's *overall* average knowledge can be represented as the number of known attributes divided by the total number of known values (ie. non Unobserved values) Figure 3 shows that overall knowledge significantly increases as the participant progresses through the simulation (t_{278} = 2.198, p = .0288). This validates that the simulation can teach basic use of force principles (assuming a correct operationalization of the knowledge attributes).

Discussion

The evidence we collected supports both of our hypotheses. Virtual reality hardware increased a subject's subjective and objective indicators of presence. Two measures of performance increased as subjects repeated the simulation.

Some subjects reported difficulty using the virtual reality controls. Multiple subjects reported that they were afraid to walk backwards for fear of bumping into walls. Walking backward was the most common way for subjects to take cover, and taking cover was required to reach the maximum score, so this may have limited some subjects' performance and may help to explain why overall score on the knowledge attribute was generally higher when using the screen and keyboard controls. These difficulties may be mitigated as virtual reality hardware becomes more common. Alternatively, we may extend the tutorial to include walking backwards so that future subjects will know it is safe.

Future Work

The first version of this simulation was tested with civilians, so having demonstrated some success, we plan to test with actual officers in training. We also plan to significantly expand the content of the simulation to include more use of force scenarios based on the feedback of experts and research by PERF (2012).

In future versions of this simulation, the only content provided by a human author will be the domain description and descriptions of possible wrong beliefs about that domain that the simulation should target during training. Wrong beliefs represent a different version of the domain. For example, an officer may not realize that approaching the agitated suspect will cause him to get angry. This misunderstanding can be represented as a version of the domain where the *get angry* axiom either does not exist or has different preconditions. The set of things a person would do differs based on their beliefs about the domain, and these differences can be used to diagnose what the trainee knows and does not know. This is how we derived the knowledge attributes used to measure subject performance—by identifying actions that only a person who knows or does not know that information would do.

The domain for this prototype was small enough that a story graph could be generated entirely offline. However, as the content of the simulation expands in size and complexity, the planner will need to generate the story graph online and automatically derive knowledge attributes. By reasoning about beliefs, the experience manager can use this model of the participant's knowledge to influence NPC actions, while still ensuring each agent acts believably, and while ensuring that only trainees who understand the knowledge attributes being taught score highly.

Acknowledgements

References

Alexander, A. L.; Brunyé, T.; Sidman, J.; and Weil, S. A. 2005. From gaming to training: A review of studies on fidelity, immersion, presence, and buy-in and their effects on transfer in pc-based simulations and games. *DARWARS Training Impact Group* 5:1–14.

Fikes, R. E., and Nilsson, N. J. 1972. STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2(3):189–208.

Hays, R. T., and Singer, M. J. 1989. *Simulation fidelity in training system design: Bridging the gap between reality and training*. Springer-Verlag New York Inc.

Ma, W.; Adesope, O. O.; Nesbit, J. C.; and Liu, Q. 2014. Intelligent tutoring systems and learning outcomes: a meta-analysis. *Journal of Edu. Psychology* 106(4):901–918.

Meehan, M. 2001. *Physiological reaction as an objective measure of presence in virtual environments*. Ph.D. Dissertation, University of North Carolina at Chapel Hill.

Newell, A.; Shaw, J. C.; and Simon, H. A. 1959. Report on a general problem-solving program. In *Proc. of ICIP*, 256–264.

Niehaus, J. M.; Li, B.; and Riedl, M. O. 2011. Automated scenario adaptation in support of intelligent tutoring systems. In *Proc. of FLAIRS*.

Police Executive Research Forum. 2012. An integrated approach to de-escalation and minimizing use of force. In *Critical Issues in Policing Series*.

Riedl, M. O., and Bulitko, V. 2013. Interactive narrative: an intelligent systems approach. *AI Magazine* 34(1):67–77.

Riedl, M. O., and Young, R. M. 2010. Narrative planning: balancing plot and character. *JAIR* 39(1):217–268.

Riedl, M. O.; Stern, A.; Dini, D.; and Alderman, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *Intl. Trans. on Systems Science and Applications Special Issue on Agent Based Systems for Human Learning* 4(2):23–42.

Riva, G.; Davide, F.; and IJsselsteijn, W. A. 2003. *Being there: Concepts, effects and measurements of user presence in synthetic environments*. Ios Press.

Roberts, D. L., and Isbell, C. L. 2008. A survey and qualitative analysis of recent advances in drama management. *Intl. Trans. on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2):61–75.

Schank, R. C.; Berman, T. R.; and Macpherson, K. A. 1999. Learning by doing. *Instructional-Design Theories and Models: A New Paradigm of Instructional Theory* 2:161–181.

Slater, M., and Steed, A. 2000. A virtual presence counter. *Presence: Teleoperators and Virtual Environments* 9(5):413–434.

Stevens, J. A., and Kincaid, J. P. 2015. The relationship between presence and performance in virtual simulation training. *Open Journal of Modelling and Simulation* 3:41–48.

Thomas, J. M., and Young, R. M. 2010. Annie: Automated generation of adaptive learner guidance for fun serious games. *IEEE Trans. on Learning Tech.* 3(4):329–343.

VanLehn, K. 2006. The behavior of tutoring systems. *Intl. Journal of AI in Education* 16(3):227–265.

Ware, S. G. 2014. Glaive: a state-space narrative planner supporting intentionality and conflict. In *Proc. of AIIDE*, 80–86.

Witmer, B. G., and Singer, M. J. 1998. Measuring presence in virtual environments: a presence questionnaire. *Presence: Teleoperators and Virtual Environments* 7(3):225–240.

Young, R. M.; Ware, S. G.; Cassell, B. A.; and Robertson, J. 2013. Plans and planning in narrative generation: a review of plan-based approaches to the generation of story, discourse and interactivity in narratives. *Sprache und Datenverarbeitung, Special Issue on Formal and Computational Models of Narrative* 37(1-2):41–64.

Zook, A.; Lee-Urban, S.; Riedl, M. O.; Holden, H. K.; Sottilare, R. A.; and Brawner, K. W. 2012. Automated scenario generation: toward tailored and optimized military training in virtual environments. In *Proc. of FDG*, 164–171.