Toward Automated Scenario Generation with Deep Reinforcement Learning in GIFT

Jonathan Rowe[1], Andy Smith[1], Robert Pokorny[2], Bradford Mott[2], and James Lester[2]
North Carolina State University[1], Intelligent Automation, Inc.[2]

INTRODUCTION

Simulation-based training is an important tool for preparing learners to perform a broad range of complex tasks and skills. A key functionality of simulation-based training environments is delivering scenarios that drive learning interactions that approximate real-world situations. However, simulation-based training scenarios are typically resource-intensive to create. Some simulation environments provide authoring tools that enable new training scenarios to be manually created by subject matter experts, but these tools often require a high degree of specialized knowledge to be utilized effectively. Authored scenarios often cannot be reused in other training environments, and the knowledge associated with particular authoring tools has limited transferability. Further, learners are usually limited to training with a finite set of training scenarios provided by the system's designers. If learners have mastered the learning objectives associated with the available set of training scenarios, there is marginal benefit provided by further training with the simulation. Finally, training simulation scenarios are often delivered following a one-size-fits-all approach: they have limited capacity to dynamically respond to the broad range of individual differences in knowledge or behavior that are typical among learners.

Automated scenario generation offers considerable promise for addressing the needs of simulation-based training. By utilizing automated scenario generation techniques, simulation environments can account for individual differences in how learners respond to different types of scenario events. Further, they can create effective variations on training scenarios without requiring every scenario to be manually authored or managed by human experts. By leveraging generative techniques from interactive narrative technologies, we can dynamically create training scenarios that are configured to address instructors' learning objectives and tailored to the cognitive and behavioral characteristics of individual students (Riedl & Bulitko, 2012; Wang et al., 2017).

Recent advances in machine learning, including artificial neural networks (in general) and deep learning (in particular), have spurred growing interest in data-driven approaches to interactive narrative generation. For example, deep reinforcement learning (deep RL) has begun to show significant promise for personalizing events in narrative-centered learning environments (Wang et al., 2017). However, there are many open questions regarding how we can most effectively leverage machine learning in order to automatically generate training scenarios that are tailored to instructors' and trainees' learning objectives. To begin to address these questions, we are launching a new collaborative effort between North Carolina State University, Intelligent Automation, Inc., and the U.S. Army Research Laboratory to investigate the design and development of a deep RL framework for automated scenario generation in GIFT. To serve as a testbed environment, we are generating training scenarios for Virtual Battlespace 3 (VBS3), a widely used simulation platform for small unit training within the Army, with an initial focus on Call for Fire (CFF) training.

In this paper, we provide an overview of the deep RL framework for automatic scenario generation. We describe how to formalize automatic scenario generation as a deep RL task. We discuss several key components of the framework, including the scenario adaptation library, simulated learners, and a deep neural network model of multi-objective rewards. We describe the VBS3 training simulation that we are utilizing as an initial testbed environment. Next, we describe preliminary findings from a proof-of-concept implementation of a reinforcement learning-based scenario generator that centers on generating initial scenario conditions for CFF training using multi-armed bandits (i.e., a stochastic scheduling technique that is related to deep RL).

AUTOMATED SCENARIO GENERATION WITH DEEP REINFORCEMENT LEARNING

We approach the task of automated scenario generation from the perspective of interactive narrative technologies. Automated scenario generation and interactive narrative generation share several key characteristics. First, in both automated scenario generation and interactive narrative generation, users are active participants in virtual worlds that dynamically respond to users' actions. Second, both automated scenario generation and interactive narrative generation center on generating sequences of events that achieve author-specified objectives to produce scenarios that are effective and engaging. Third, both automated scenario generation and interactive narrative generation produce scenarios that are realized in immersive simulation environments.

We formulate automated scenario generation as an instance of data-driven interactive narrative generation using deep RL (Wang et al., 2017). Deep RL is a computational framework that integrates two complementary families of machine learning techniques: reinforcement learning methods for training models for sequential decision-making under uncertainty, and deep neural networks for pattern recognition and representation learning with big data. Reinforcement learning is the task of a software agent inducing a control policy for selecting actions in an uncertain environment with delayed rewards (Sutton & Barto, 1998). Deep neural networks combine weighted summations of non-linear functions to extract and model multi-layer hierarchical representations of data using supervised, semi-supervised, and unsupervised machine learning techniques (Goodfellow, Bengio, & Courville, 2016). By integrating reinforcement learning and deep neural networks, deep RL provides a formalized framework for sequential decision-making in complex environments.

Deep reinforcement learning provides a natural computational framework for formalizing dynamic scenario generation: the generator is tasked with making a series of decisions about how specific scenario events should unfold at runtime to optimize student performance on a pre-specified set of learning objectives. Dynamic scenario generation can be modeled as a sequential decision-making task in which a scenario generator introduces successive adaptations to scenario events over discrete time steps. A time step represents the time point when an adaptable event, such as the introduction of an obstacle or elimination of a resource, is triggered in the scenario. Using this formalization, deep RL can be utilized to dynamically generate adaptive “child” training scenarios from a canonical (i.e., “parent”) scenario that explicitly optimizes for both author-specified objectives and trainee learning outcomes. By inducing multi-objective reward models for controlling run-time decisions about training scenario events, we intend to enable authors to specify learning objectives that generate personalized training scenarios in immersive simulation environments integrated with GIFT.

The deep RL framework for automated scenario generation consists of several key components: (1) a deep Q-Network model for controlling run-time scenario adaptation decisions that optimize multiple scenario objectives, (2) a scenario adaptation library that specifies possible transformations of “parent” scenarios to generate “child” scenarios, and (3) a simulated learner framework for generating synthetic data to train an initial version of the scenario generator. In addition, the framework requires a software infrastructure for integrating automated scenario generation functionalities with GIFT’s modular software architecture.

Deep reinforcement learning leverages a Q-Network, a type of deep neural network, to model the estimated Q values of state-action pairs gathered from past observations of student interactions with a scenario during reinforcement learning. Q-networks encode the expected benefits of specific scenario adaptations in terms of a "reward function," an explicit mathematical expression of optimization criteria that guide automated scenario generation. In the original work on deep reinforcement learning for Atari game playing, the Q-network was organized as a convolutional neural network, which is a natural choice for processing image data from 2D games. For automated scenario generation, we will investigate deep architectures that utilize ***long-short term memory networks*** (LSTMs), a type of recurrent neural network, for modeling sequential data as typically expected in simulation scenarios. LSTMs are specifically designed for processing sequences of temporal data. LSTMs have achieved high predictive performance in many sequence labeling tasks, often outperforming standard recurrent neural networks by using a longer-term memory than standard RNNs, preserving short-term lag capabilities, and effectively addressing the vanishing gradient problem. We anticipate that utilizing LSTMs will enable reinforcement learning-based scenario generators to extract complex nonlinear interaction patterns between observed events and scenario adaptation decisions. LSTMs will be utilized to implement multi-objective deep Q-networks for automated scenario generation, as well as to implement machine learning-based simulated students to generate synthetic training data in future work.

Multi-objective reward functions will enable the automated scenario generator to consider tradeoffs between competing authorial goals, learning objectives, and learner engagement. This builds upon prior research by the NCSU team on multi-objective RL for interactive narrative generation (Sawyer, Rowe, & Lester, 2017), and it involves incorporating a vector-based representation of reward in the output layer of a deep Q-network, where vector indices correspond to different reward sources. A multi-objective Q-Network is induced at training time, and it yields a run-time scenario adaptation model after the course author has specified relative preferences among competing reward sources at course creation-time.

State representations for driving deep RL-based scenario generation decisions will consist of several complementary sources of information. First, state vectors will include domain-independent features encoding learner knowledge, traits, and performance characteristics. A key requirement for automated scenario generation is devising generalized assessment rules that can be applied to a broad range of generated scenarios within a given task domain; it would be prohibitive for a system designer to devise custom assessment logic for every automatically generated scenario. Second, state vectors will include several features that summarize the history of past scenario adaptation decisions performed by the scenario generator. Third, state vectors will include a one-hot encoding of the category of scenario adaptation decision under consideration in order to leverage modularity and maintain tractability of the reinforcement learning task. These state features are consistent with the Adaptive Tutoring Learning Effect Chain, and they are consistent with our project's vision for investigating how scenario generation functionalities should most effectively be integrated with the Pedagogical and Domain Modules of GIFT.

Scenario Adaptation Library

A key component of the deep RL-based scenario generation framework is devising a scenario adaptation library, which enumerates the range of possible transformations to a "parent" scenario that can be applied to generate "child" scenarios. By investigating different combinations of prospective scenario adaptations, the deep RL framework can generate a vast range of possible training scenarios that can be deployed with simulated or human learners, evaluated for their effectiveness in terms of trainee learning outcomes, and utilized to refine the scenario adaptation model for adaptive personalized scenario generation.

Integrating deep RL-based scenario generation into GIFT is a key aspect of the project. A key interest is exploring potential extensions to GIFT that support domain-independent specifications of scenario adaptations—these would be specified with GIFT's Pedagogical Module—in line with project objectives of generalizability of deep RL-based scenario generation. We envision a generalized taxonomy of scenario adaptations that includes several hierarchical domain-independent categories, including (1) inserting or removing obstacles; (2) constraining or increasing resources; (3) reconfiguring key objects; (4) adding, modifying, or removing sub-tasks; and (5) providing or removing embedded scaffolding. These categories will characterize a range of candidate adaptations that can be applied to a parent scenario in order to generate a set of “child” scenarios. These domain-independent scenario adaptations could be instantiated within GIFT’s Domain Module, which will configure and launch scenario events at runtime via a GIFT gateway to be realized in the simulation-based virtual training environment.

Simulated Learners

In order to train deep reinforcement learning-based models of dynamic scenario generation, we will utilize synthetic training data produced by simulated students created for each of the virtual training environments. The design of simulated students is informed by related work in artificial intelligence in education (McCalla & Champaign, 2013) and spoken dialog systems (Schatzmann, Weilhammer, Stuttle, & Young, 2006). We investigate how simulation parameters related to model granularity and model complexity influence synthetic data generation for deep reinforcement learning-based scenario generation (Rowe et al., 2017).

Dynamic Scenario Generation User Experience

The user experience of automated adaptive scenario generation functionalities in GIFT is likely to be different based on whether the user is a course designer, a student, or a software developer. For a pre-integrated training environment, a course designer will select the training objectives that he is targeting in the GIFT Course Creator, and he can specify constraints on specific scenario adaptations that he would like to avoid in the generated run-time scenario. As long as the deep RL-trained scenario generator can produce a scenario that is consistent with the objectives and constraints provided by the author, the course will validate and it can be tested with live students. For a student, automated scenario generation will be invisible, and training events will be tailored based on the student’s individual traits, knowledge, and performance in the simulation environment.

For a software developer seeking to integrate a new domain or training application, she will need to (1) have a deep knowledge of the “parent” scenarios supported by the training environment, (2) create a specification of possible “child” scenario adaptations that can be realized in the training environment, (3) develop a gateway module that mediates communication between scenario generation functionalities in GIFT and the training application, (4) have access to training data for inducing deep reinforcement learning-based scenario generation models if existing domain-independent models cannot be reused, and (5) integrate trained scenario generation models into run-time GIFT courses. Given these resources, a software developer will be able to use deep reinforcement learning-based scenario generation functionalities to create a new scenario generator for a novel domain or simulation environment.

GIFT Integration

Integrating deep RL-based scenario generation into GIFT is a key objective of the project, and supporting automated scenario generation has several implications for the GIFT architecture, authoring tools, and software. For example, supporting automated scenario generation in GIFT will likely involve extensions to the GIFT Course Creator to enable instructors to identify relevant learning objectives that should guide the generation of relevant training scenarios. Further, devising tools for ranking and visualizing automatically generated training scenarios will be essential for instructors to configure scenario generation functionalities for use in training courses that they create. Devising generalized assessment logic that can operate across multiple scenarios, and be specified in GIFT DKF files, will be critical for ensuring that course creators do not need to hand-specify custom assessment rules for every generated scenario. Finally, the project seeks to investigate support for domain-independent specifications of scenario adaptations—these would be specified by GIFT's Pedagogical Module—in line with project objectives of generalizability of deep RL-based scenario generation. This generalized taxonomy of scenario adaptations will include hierarchical domain-independent categories, such as (1) inserting or removing obstacles; (2) constraining or increasing resources; (3) reconfiguring key objects; (4) adding, modifying, or removing sub-tasks; and (5) providing or removing embedded scaffolding. These categories characterize a scenario adaptation library that defines the space of possible generatable scenarios in a manner that holds potential for portability and reuse.

Figure 1. Screenshot of Virtual Battlespace 3 simulation environment.

VIRTUAL BATTLESPACE 3 TESTBED ENVIRONMENT

The selection of testbed simulation-based virtual training environments is guided by two key requirements. First, the simulation environment should either be open-source, or include APIs or tools for generating novel training scenarios, as well as models for specifying adaptations to scenarios at run-time. By enabling close integration between GIFT and the simulation-based training environment, it is possible to engage in rapid iteration cycles for designing, developing, and testing directions for dynamic scenario generation. Second, the selection of the simulation-based virtual training environment testbeds should prioritize environments that support scenario generation that enable the scenario generator to produce thousands (or more) of “child” scenarios from a single “parent” scenario. To fully exercise the deep RL framework’s generative capabilities (i.e., its ability to broadly explore a given scenario space) and fully stress test its computational capabilities, testbeds should include a broad range of event types, actors, and trainee interactions.

The primary testbed simulation environment for the first phase of this project will be Virtual Battlespace 3 (VBS3). Built by Bohemia Interactive Simulations, VBS3 is the Army’s most widely used simulation platform for small unit training (Figure 1). Designed as a flexible simulation tool for tactics training and mission rehearsal, VBS3 provides realistic physics, high-fidelity 3D graphics, expansive geo-specific terrains, and a large content library of 3D digital assets. VBS3 can be used for a broad range of training purposes, including training for cordon and search of specific structures, breaching obstacles, defense of territory with machine gun and mortars, and clearing highways of IEDs. VBS3 also includes features that enable dynamic modifications to training scenarios, as well as features for observation of the environment by instructors, and an After Action Review playback capability.

Although it is a closed-source training simulation, VBS3 provides several developer tools that can facilitate research on automated scenario generation, including a real-time scenario editor, an offline mission editor, tools for importing new 3D assets, and flexible terrain creation functionalities. VBS3 is used widely in the U.S. Army, and it is integrated with GIFT 2017-1 through a previously developed gateway module. Further, our work with VBS3 will build upon prior research by IAI to devise low-cost assessment frameworks for intelligent tutoring systems through feedback from subject matter experts.

During the first year of the project, we will focus on automated scenario generation in the task domain of Call for Fire training. The CFF task domain in VBS3 will encompass scenarios in which an infantry soldier requests indirect fire from supporting artillery (e.g., unmanned aircraft) on an identified target. The steps of this task include identifying the target, waiting for the artillery to move into position, calling for artillery fire using an established communication protocol, adjusting artillery fire, and providing a damage assessment. "Child" scenarios in the Call for Fire task domain will modify the type, visibility, and movement of the target; augment surrounding terrain and vegetation; change the weather and time-of-day; impact radio communications with artillery operators; augment the type of artillery fire (e.g., smoke, explosive); and influence the accuracy and damage of the artillery fire.

PRELIMINARY FINDINGS ON AUTOMATED GENERATION OF INITIAL SCENARIO CONDITIONS

As a starting point, we developed a prototype system using a multi-armed bandit (MAB) computational formalism, consisting of several components of the proposed deep RL pipeline. The MAB implementation utilizes initial versions of a scenario adaptation library, a simulated learner, and a reward function.

A multi-armed bandit is a class of sequential decision problem in which a set of resources must be allocated between competing choices. MABs are related to reinforcement learning, but they do not account for stochastic effects of sequential decisions on environment states. In an MAB, each choice, or arm, has a defined reward unknown to the system, thus requiring it to explore different choices to learn which of the choices provides optimal expected reward over a finite series of trials. Typically, bandit algorithms are designed to minimize regret, which is the difference between the reward accumulated by the system and the reward the system would have received if it had pulled the optimal arm at every trial. Depending on a variety of factors such as the type of rewards (stochastic, non-stochastic) or the type of regret being minimized (instantaneous or cumulative), different algorithms have been shown to obtain near optimal solutions (Vermorel, 2005). Variants of MABs have been shown to be an effective solution for a variety of tasks such as sequencing learning activities (Liu, 2014) and playing real-time strategy games (Ontanón, 2017).

The first step in formalizing scenario generation as a MAB is defining a scenario adaptation library for the Call for Fires task. As MABs do not have a concept of state, and thus do not capture changes in the simulation environment, we focus our scenario adaptation library on initial conditions of Call for Fires scenarios. In this prototype we focus on 3 categories of initial conditions: weather, time of day, and target mobility. These categories were chosen because of they affect the difficulty of a Call for Fires scenario, and they can also be realized in the VBS3 environment We defined three possible values for weather (clear, cloudy, rain), three possible values for time of day (day, dusk, night), and two possible values for target mobility (still, moving). This corresponds to 18 possible scenarios that could be generated and evaluated by the MAB.

Next, to provide data to train the system we created a set simulated learners. Each simulated learner consists of a competency score from 0 to 1, representing their ability for a Call for Fire task. To generate rewards for each scenario, we crafted a reward function that takes into account both the difficulty of the scenario and the skill level of the student. Difficulty levels were authored for each type and value of initial condition, with values being averaged to determine the difficulty of each generated scenario. The difficulty level was then combined with the learner's competency score to generate the probability that the learner would increase their competency level from creating the exercise. A 0 or 1 reward was then generated for the trial by sampling from a Bernoulli distribution that was parameterized using the combination of scenario difficulty and learner competency level.

We ran MAB simulations for two populations of simulated learners. For the first simulation, a learner was selected for each trial from a Gaussian distribution centered around a "low" competency score (*M* = .2, *SD* =.1). For the second simulation, learners were selected from a distribution centered around a high competency score (*M* = .8, *SD* = .1). For each simulation, we ran 50,000 trials of an 18-armed bandit using the UCB1 algorithm to manage exploitation/exploration of the arms. Figure 2 shows the average rewards of the top-5 arms (i.e., generated scenarios) for both types of simulated learners. We observe that after some shuffling, each arm begins to stabilize around the "true" reward for that given scenario. For the Low Competency learner group, the scenarios recommended are all "easier" scenarios with non-moving targets and high visibility, which is to be expected given that our reward formulation does not expect low-competency learners to benefit significantly from difficult scenarios. Similarly, the High Competency learner group favors more difficult scenarios featuring moving targets and poor weather/visibility.

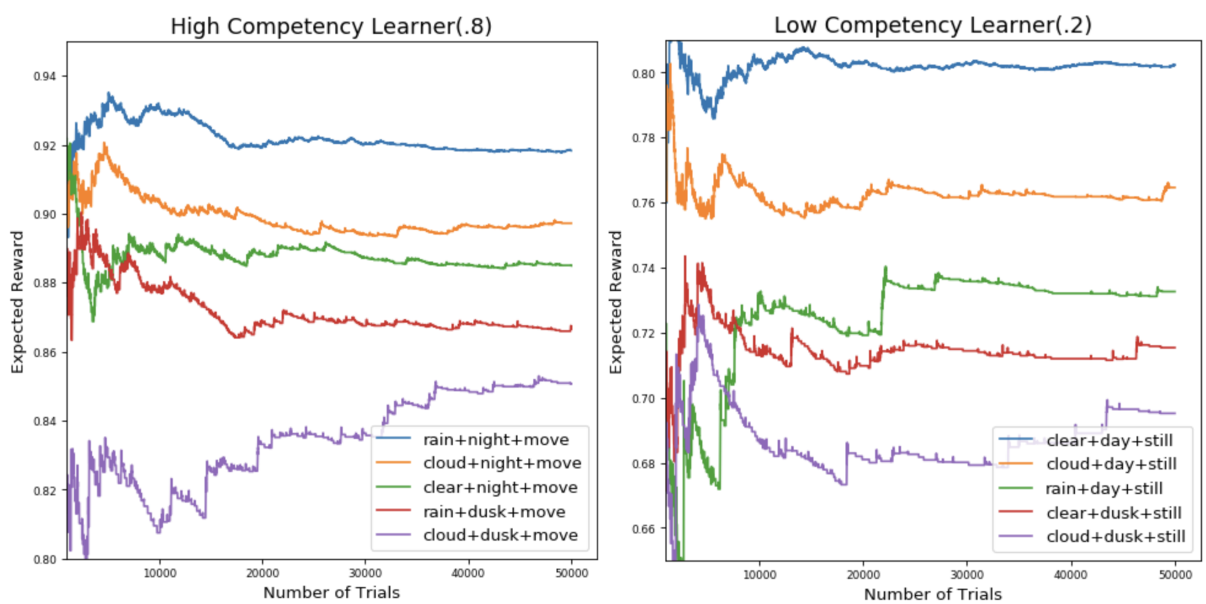

Figure 2. Average rewards over MAB trials for top-5 generated CFF scenarios for high- and low-competency simulated learners.

This prototype highlights key design considerations for deep RL-based scenario generation, but it also has several limitations. First, since MABs have no concept of state, they are not necessarily the ideal formalism for generating and evaluating dynamic, adaptive training scenarios required by more complex CFF tasks; MABs are well suited for generating the initial conditions of simple training scenarios but not sequential events. In future iterations, we will utilize reinforcement learning techniques that account for sequential decisions in order to address this additional source of complexity in scenario generation.

A second limitation is that our current simulated learner and reward models only consider one competency and reward source. As we move forward, the system will need to consider multiple learning objectives and trade-offs between them. Additionally synthesized data will eventually need to be replaced or validated with data from real human learners.

CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

Automated scenario generation is likely to serve a key role in the future of simulation-based training because of its significant potential to reduce the cost of creating novel scenarios and expand access to high-quality simulation-based training. Data-driven approaches to automated scenario generation hold promise for enhancing trainee learning experiences by leveraging recent advances in both machine learning and interactive narrative technologies. We have presented an overview of a deep RL framework for data-driven automated scenario generation, which formalizes the task in terms of sequential adaptations to a canonical "parent" scenario in order to generate "child" scenarios that can be evaluated with simulated or human learners to assess learning outcomes. Automated assessments of trainee learning outcomes drive the generator to iteratively refine its scenario generation policies and tailor scenario generation to individual learners and instructor training objectives. During the initial stages of the project, we are investigating deep RL-based scenario generation in the context of CFF training using the VBS3 simulation environment. To serve as an initial proof-of-concept for data-driven automated scenario generation, we conducted a preliminary investigation of multi-armed bandit techniques for generating initial conditions of CFF training scenarios. Preliminary results indicate that multi-armed bandits, combined with a simple simulated learner model and scenario adaptation library, can produce a ranked ordering of automatically generated training scenarios that are tailored to learners' individual differences.

In future work, we plan to significantly expand the scenario adaptation library to capture a broader range of possible transformations to "parent" training scenarios, including sequential adaptations that can be performed dynamically over the course of a scenario. This will allow us to expand our formulation of automated scenario generation beyond initial scenario conditions and begin exploring deep RL techniques. Further, we plan to investigate richer simulated learner models that can serve as a bootstrapping mechanism for automated scenario generation, as well as multi-objective rewards to enable scenario generation that accounts for complex tradeoffs between complementary and competing learning objectives. Finally, we plan to investigate manual, semi-automated, and automated techniques for realizing generated scenarios in VBS3, enabling human learners to interact with adaptive training scenarios that have been generated using deep RL.

ACKNOWLEDGMENTS

The research described herein has been sponsored by the U.S. Army Research Laboratory under cooperative agreement W911NF-18-2-0020. The statements and opinions expressed in this article do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

REFERENCES

Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep Learning (Vol. 1)*. Cambridge: MIT press.

Liu, Y. E., Mandel, T., Brunskill, E., & Popovic, Z. (2014). Trading Off Scientific Knowledge and User Learning with Multi-Armed Bandits. *Proceedings of the 7th International Conference on Educational Data Mining* (EDM-2014), London, UK, pp. 161-168.

McCalla, G. & Champaign, J. (2013). Simulated Learners. *IEEE Intelligent Systems*, 28(4), 67-71.

Ontanón, S. (2017). Combinatorial multi-armed bandits for real-time strategy games. *Journal of Artificial Intelligence Research*, 58, 665-702.

Riedl, M. O., & Bulitko, V. (2012). Interactive narrative: An intelligent systems approach. *AI Magazine*, 34(1), 67.

Rowe, J., Pokorny, B., Goldberg, B., Mott, B., and Lester, J. (2017). Toward Simulated Students for Reinforcement Learning-Driven Tutorial Planning in GIFT. *Proceedings of the 5th Annual GIFT User Symposium (GIFTSym5)*. Orlando, Florida

Sawyer, R., Rowe, J., & Lester, J. (2017). Balancing Learning and Engagement in Game-Based Learning Environments with Multi-Objective Reinforcement Learning. *Proceedings of the 18th International Conference on Artificial Intelligence in Education (AIED-2017)*, Wuhan, China, pp. 323-334.

Schatzmann, J., Weilhammer, K., Stuttle, M., & Young, S. (2006). A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *The Knowledge Engineering Review*, 21(2), 97-126.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Vermorel, J., & Mohri, M. (2005). Multi-Armed Bandit Algorithms and Empirical Evaluation. *Proceedings of the 16th European Conference on Machine Learning (ECML-2005)*, Porto, Portugal, pp. 437-448.

Wang, P., Rowe, J., Min, W., Mott, B., & Lester, J. (2017). Interactive Narrative Personalization with Deep Reinforcement Learning. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-2017), Melbourne, Australia, pp. 3852-3858.

ABOUT THE AUTHORS

Dr. Jonathan Rowe *is a Research Scientist in the Center for Educational Informatics at North Carolina State University. He received the Ph.D. and M.S. degrees in Computer Science from North Carolina State University, and the B.S. degree in Computer Science from Lafayette College. His research is in the areas of artificial intelligence and human-computer interaction for advanced learning technologies, with an emphasis on game-based learning environments, intelligent tutoring systems, user modeling, educational data mining, and computational models of interactive narrative generation.*

Mr. Andy Smith *is a Research Scientist in the Center for Educational Informatics at North Carolina State University. He received his M.S. in Computer Science from North Carolina State University, and his B.S. degrees in Computer Science and Electrical and Computer engineering from Duke University. Prior to graduate school Andy worked as an Underwater Robotics Engineer at SPAWAR SSC Pacific in San Diego, CA. His research is focused on the intersection of artificial intelligence and education, with emphasis on user modeling, game-based learning, and educational data mining.*

Dr. Robert Pokorny *is Principal of the Education and Training Technologies Division at Intelligent Automation, Inc. He earned his Ph.D. in Experimental Psychology at University of Oregon in 1985, and completed a postdoctoral appointment at University of Texas at Austin in Artificial Intelligence. Bob's first position after completing graduate school was at the Air Force Research Laboratory, where he developed methodologies to efficiently create intelligent tutoring systems for a wide variety of Air Force jobs. At Intelligent Automation, Bob has led many cognitive science projects, including adaptive visualization training for equipment maintainers, and an expert system approach for scoring trainee performance in complex simulations.*

Dr. Bradford Mott *is a Senior Research Scientist in the Center for Educational Informatics at North Carolina State University. Prior to joining North Carolina State University, he served as Technical Director at Emergent Game*

Technologies where he created cross-platform middleware solutions for video game development, including solutions for the PlayStation 3, Wii, and Xbox 360. Dr. Mott received his Ph.D. in Computer Science from North Carolina State University in 2006, where his research focused on intelligent game-based learning environments. His current research interests include computer games, computational models of interactive narrative, and intelligent game-based learning environments.

Dr. James Lester *is Distinguished Professor of Computer Science at North Carolina State University, where he is Director of the Center for Educational Informatics. His research centers on transforming education with technology-rich learning environments. With a focus on adaptive learning technologies, his research spans intelligent tutoring systems, game-based learning environments, affective computing, and tutorial dialogue. The adaptive learning environments he and his colleagues develop have been used by thousands of students in K-12 classrooms. He received his Ph.D. in Computer Science from the University of Texas at Austin in 1994. He is a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI).*