

Reflective Tutoring for Immersive Simulation

H. Chad Lane, Mark Core, Dave Gomboc, Steve Solomon, Michael van Lent,
and Milton Rosenberg

Institute for Creative Technologies, University of Southern California
13274 Fiji Way, Marina del Rey, CA, 90292 USA
{lane,core,gomboc,solomon,vanlent,rosenberg}@ict.usc.edu

Abstract. Reflection is critically important for time-constrained training simulations that do not permit extensive tutor-student interactions during an exercise. Here, we describe a reflective tutoring system for a virtual human simulation of negotiation. The tutor helps students review their exercise, elicits where and how they could have done better, and uses explainable artificial intelligence (XAI) to allow students the chance to ask questions about the virtual human’s behavior.

1 Introduction

Reflection is widely regarded as a critical meta-cognitive skill for learning. Not surprisingly, researchers have found that human tutors often continue discussing a problem with a student after it has been solved [1]. For problem solving that occurs under hard time constraints (e.g., flight training), reflection is even more important since there is limited time for decision making during practice. Additionally, the tutor must avoid becoming a distraction (e.g., by giving too much advice). In business and the military, it has long been recognized that the *after-action review* (AAR) is when a bulk of the learning occurs [2].

In this paper, we report on an intelligent tutoring system (ITS) that scaffolds reflection activities with the student, such as reviewing salient events from an exercise, discussing ways to improve, and asking questions of entities involved in the simulation. Salient events typically include critical decision points in the simulation, errors of omission and commission, and also instances of positive actions on the part of the student. Question-asking is achieved by the use of explainable artificial intelligence (XAI) [3], a simulation-independent approach to explanation with roots in earlier simulation-dependent explanation research (e.g., [4]). XAI allows entities to describe their reasoning, choices, and state changes through questions derived from their AI models. Using XAI as a standalone tool would be tantamount to unguided discovery [5], and so our ITS monitors its use and gives hints on the best questions to ask. This differentiates our approach from other reflective systems, such as SCoT-DC [6].

2 Planning and Executing a Reflective Tutoring Session

To conduct a reflective tutoring session, it is necessary to both plan ahead and execute the plan. Because the student’s answers might reveal new misconcep-

1	T	In the simulation history window, can you select a line where you could have done better?
2	S	Line 5
3	T	Why do you think line 5 upset the doctor?
4	S	He didn't want to move his clinic
5	T	That's right. You abruptly admitted your goal was to have him move his clinic. What else could you have said?
6	S	I don't know
7	T	You needed to build more familiarity. For example, you could have asked how he was doing today. Let's focus on line 10.
8	T	In the XAI window, please ask the doctor some questions and find out how this utterance contributed to failure of the negotiation.
9	S	Why did the negotiation fail?
10	D	I no longer trusted the captain.
11	S	How was your solidarity with the Captain affected here?
12	D	It decreased.
13	S	Why did your solidarity with the Captain decrease?
14	D	I disagreed with the intentions of the Captain expressed in this utterance.
15	T	Good job. In general, you want to avoid revealing sensitive information unless absolutely necessary.

Fig. 1. A reflective dialogue in which the student learns that it was unwise to reveal an undesirable fact before sufficiently building trust. References to “lines” are to the actual exercise (not shown here) and questions asked in the XAI subdialogue were selected from a menu.

tions, it is important that the tutor be able to adapt an AAR plan on the fly. Our system begins its planning process by loading a log file from the target simulation and performing the following steps:

1. **analyze student's exercise:** highlight important events from the exercise that are candidates for discussion.
2. **create agenda:** organize and prioritize the highlighted events.
3. **prepare XAI:** load exercise log, action representations, and natural language generation knowledge (details in [3]).

The first two steps roughly model what human instructors need to do to perform an AAR: judge the student's performance, make decisions about what merits discussion, and finally, decide how they might go about addressing these issues. Currently, steps 1 and 2 require human support, but we are working on automating these tasks as part of an in-game tutor that assesses turn-by-turn choices of the student. The resulting agenda is then passed to a planner and executor that conduct the dialogue – an example appears in figure 1. Prior to this, the student had completed a session with a virtual doctor who is running a clinic in a dangerous location [7]. The student's task is to convince the doctor to move willingly to a safer location through building trust and bargaining.

The reflective tutor's actions are determined by a hierarchical task network planner. Our prototype uses 12 recipes that implement various reflective activities, such as asking the student to identify mistakes (e.g., line 1 of the figure), suggesting ways to improve (line 7), and using XAI to perform "investigations" (lines 9-14). To support XAI, we use a simple model of investigation comprised of a sequence of ideal questions and associated hints that are given when the student fails to ask the right questions. Natural language generation is accomplished via templates and we currently use a keyword-based approach to handling answers to open-ended questions (e.g., line 4).

3 Ongoing and Future Work

We are currently porting our system to a serious game for teaching cultural awareness and negotiation. Although the version of our system presented here assumes no tutor presence during an exercise, our new version coordinates the reflective activities with advice received during the simulation (Katz et. al. refer to this as *distributed* tutoring [1]). We are also exploring more advanced natural language generation and understanding techniques.

Acknowledgments

Special thanks to Bill Swartout and David Traum for their guidance and support. The project described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the U.S. Government, and no official endorsement should be inferred.

References

1. Katz, S., Allbritton, D., Connelly, J.: Going beyond the problem given: How human tutors use post-solution discussions to support transfer. *International Journal of Artificial Intelligence in Education* **13** (2003) 79–116
2. Morrison, J.E., Meliza, L.L.: Foundations of the after action review process. Technical Report 42, U.S. Army Research Institute (1999)
3. Core, M.G., Lane, H.C., van Lent, M., Gomboc, D., Solomon, S., Rosenberg, M.: Building explainable artificial intelligence systems. In: Proc. of the 18th Conference on Innovative Applications of Artificial Intelligence (IAAI06), Boston, MA (in press)
4. Johnson, W.L.: Agents that learn to explain themselves. In: Proceedings of the Twelfth National Conference on Artificial Intelligence. (1994) 1257–1263
5. Mayer, R.: Should there be a three-strikes rule against pure discovery learning? the case for guided methods of instruction. *American Psychologist* **59**(1) (2004) 14–19
6. Peters, S., Bratt, E.O., Clark, B., Bon-Parry, H., Schultz, K.: Intelligent systems for training damage control assistants. In: Proc. of Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC04), Orlando, FL (2004)
7. Traum, D., Swartout, W., Marsella, S., Gratch, J.: Fight, flight or negotiate: Believable strategies for conversing under crisis. In: Proc. of the 5th International Working Conference on Intelligent Virtual Agents. (2005)