THESPIAN: An Architecture for Interactive Pedagogical Drama

Mei Si, Stacy C. Marsella and David V. Pynadath
Center for Advanced Research in Technology for Education
USC Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA, 90292 USA
{mei,marsella,pynadath}@isi.edu

Abstract. Interactive drama is increasingly being used as a pedagogical tool in a wide variety of computer-based learning environments. However, the effort required to build interactive dramas is quite significant. We built Thespian, an architecture that supports faster development of IPDs, open-ended interaction, encoding of pedagogical goals and quantitative metrics for evaluating those goals. Thespian uses autonomous agents to control each character and assumes that the starting point for the design process is a set of standard scripts. A "fitting" algorithm facilitates the design process by automatically adjusting the goals of the agents so that the agents perform their roles according to the scripts. This also ensures the agents will behave true to their character's motivations even when the interactive drama deviates from the scripts. In this paper, we discuss this basic approach in detail and illustrate its application to the Tactical Language Training System.

Keywords. pedagogical agents, authoring & assessments tools, language learning

1. Introduction

Interactive drama is increasingly being used as a pedagogical tool in a wide variety of computer-based learning environments (e.g., [5,6,9,10]). In an interactive pedagogical drama (IPD), the learner interacts with the characters in a story and the story unfolds based on those interactions. Ideally, an IPD combines the pedagogical power of drama with a more active learning experience that allows learners to explore a simulated story world and see the effect of their actions. The engaging nature of drama and the direct link between actions and outcomes ideally engages students more, motivates them to spend more time learning (e.g., to explore possible paths in the story), and appropriately contextualizes the experience.

However, the creation of interactive pedagogical drama faces several challenges. Up to now, the effort required to design and build interactive dramas is quite significant [3,8], potentially requiring man-years of design and implementation. Further, effective design for relatively open-ended user interactivity is still an open research issue. And there is often a tension between the goal of interactivity and the goal of creating an engaging drama with consistent, well-motivated characters. Satisfying both goals can be a significant technological and creative challenge. For example, writers often do not have expertise in designing interactive stories, which is still largely a nascent art form. More fundamentally, in an interactive pedagogical drama, pedagogical goals must also be achieved. This raises the question of how the pedagogy is embedded in the environment,

how the experience of playing the game leads to desired learning outcomes and what the metrics are that determine whether pedagogical goals have been achieved.

We have developed an approach that speeds up the development of IPDs, supports open-ended interaction, achieves pedagogical and dramatic goals and supports quantitative metrics for evaluating the learner's achievement. We call our system Thespian, due to its actor-centric approach to realizing IPDs. Thespian's basic architecture uses autonomous software agents to control each character, with the character's personality and motivations encoded as agent goals. The ability of goal-driven agents to autonomously select actions based on the current state of the world allows them to be responsive to open-ended user interactions, while staying consistent with their "personality". We ensure that the learner's experience in the drama is consistent with pedagogical goals by embedding them in the drama; the world and characters in the world behave in ways that reinforce the lessons that the IPD is trying to teach. We can then define quantitative metrics on the achievement of pedagogical goals in terms of what happens in the story.

Thespian assumes that the starting point for the design process is a standard script or story outline, with possible variations, produced by a writer. This approach is typically used (e.g., [3]) because it provides a good baseline for creating an experience that can satisfy dramatic and pedagogical goals. The problem is that going from such linear script material to an interactive agent-based system is an arduous, time-consuming process requiring extensive software skills. We significantly facilitate the process by using an automated "fitting" algorithm [11] that adjusts agents' goals so that they are motivated to perform their roles according to the scripts. This ensures that the agents' autonomous behavior can follow the script when the learner's behavior is consistent with it, but is still true to their character's motivations even when the drama deviates from the script.

In this paper, we discuss this basic approach in detail. We also illustrate its application to the Tactical Language Training System (TLTS) [4] for rapidly teaching students the rudiments of a foreign language and culture.

2. Example Domain: Tactical Language Training System

TLTS is comprised of two main components that mutually reinforce the learning experience: a Mission Skill Builder (MSB), and a Mission Practice Environment (MPE). We will limit our discussion here to the MPE, a 3D role-playing interactive drama for learners to practice using their language skills. In the drama, the learner takes on the role of an army sergeant assigned to conduct a civil affairs mission in a foreign town. The learner navigates in the virtual world and interacts with virtual characters using spoken Arabic and gestures. The MPE contains several training scenes, each requiring the learner to carry out specific tasks within the interactive drama. We focus on the first scene to illustrate our approach to building an IPD. The story begins in a village café. The learner's mission is to establish rapport with the local people and find out about their leader. The learner enters the café and interacts with several of the locals, including an old man and

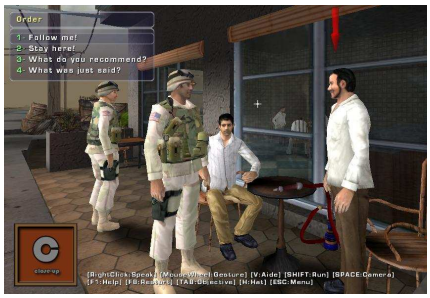

Figure 1.Scene 1 from the MPE.

a young man. The difficulty of the mission varies according to the learner's language skills. In the novice level, both of the locals are relatively cooperative, while in the expert level, the young man worries more about the safety of the town than being helpful. He may accuse the learner of being a CIA agent if he fails to establish trust. If, on the other hand, the learner uses culturally appropriate behavior, the old man will assist them.

3. Desiderata for IPDs

The design of an interactive pedagogical drama must address several requirements. Foremost, the characters should be well-motivated. They should behave according to the scripts when the learner's behavior is consistent with it. Additionally, when there are deviations from the scripts, the characters should behave consistently with the motivations implied by their on-script behavior.

Most critically, an IPD needs to encode the pedagogy. There must be some way for the learner to interact with the system and learn on the basis of that interaction. In particular, a system that has an explicit awareness of its own pedagogical goals is better equipped to work toward them in the face of unexpected interactions with the learner.

Characters must also support and maintain the interaction with the learner. Any dialog-based interactive drama must support social interaction in the form of a dialog with the learner. In practice, this means that characters should follow social norms, unless strongly motivated to violate them. By following norms, characters behave consistent with expectations and provide an incentive to interact with them.

Finally, from a drama designer's perspective, the design process should be as free of technical burdens as possible and ideally support reuse of previously developed materials.

4. Thespian

We developed Thespian as a multiagent system for controlling virtual characters in an IPD. Thespian builds on top of PsychSim, a multi-agent system [7] that controls the characters. PsychSim provides a framework for goal-driven, social behavior that forms a sound basis for meeting the requirements of IPDs that we discussed in Section 3. PsychSim agents generate their behavior through a bounded planning process that seeks to achieve their goals. Thus, the agents will choose only those behaviors that are consistent with their character profiles. PsychSim agents have a "Theory of Mind" that allows them to form mental models about other entities in the IPD, including the learner. Thus, we can potentially encode pedagogical goals as desired conditions on our model of the student. These mental models also allow a PsychSim agent to reason about the effects of its behavior on its relationships with other entities. This social reasoning capability can encode the social norms that support and maintain interactions with the user. Finally, PsychSim provides algorithms for tuning model parameters in response to the desired agent behavior. We can apply such algorithms to simplify the authoring process by ensuring that characters behave according to the script when the learner's behavior is consistent with it. This section describes how we built Thespian on top of these basic capabilities.

4.1. Goal-Driven Behavior

PsychSim represents goals as degrees of achievement with regard to certain state features (physical features, relationships, knowledge, etc.). The agents make decisions on what action to perform or what to say based on their beliefs on the possible effects of such decisions. Actions change the physical world in some fixed (possibly uncertain)

way. Saying something to another agent changes the beliefs of that agent and of any other agent that may overhear. The agents project into the future to evaluate the effect of each option on the state and beliefs of the other entities in the IPD. The agents consider not just the immediate effect, but also the expected responses of the other entities and, in turn, the effects of those responses. The agent evaluates the overall effect with respect to its goals and then chooses the action that has the highest expected value. From a decision-theoretic viewpoint, we can view this decision procedure as a boundedly rational variation on the standard solution of a *partially observable Markov decision problem (POMDP)* [13]. Thus, every action chosen by an agent is motivated by its goals, although irrational behavior may still arise due to erroneous beliefs.

We use PsychSim's basic goal representation to encode the many possible goals that our Thespian agents may have. We draw from a goal taxonomy from the psychological literature [2]. Many of these goals will conflict with each other in everyday situations. The standard "achievement" goals of logical representations are insuffi cient to resolve such conflicts because of the ambiguity that arises. PsychSim's decision-theoretic representation allows Thespian to model different character profi les by varying an explicit relative priority among the set of possible goals. Thus, Thespian models a character profi le as its various goals and their relative importance (weight). For example, in the MPE, the old man has goals of maximizing its safety level and maximizing the level of being likable, with the latter being weighted as more important. Varying these relative weights leads to changes in the agent's behavior, giving us a wide range of possible characters that will all still act in a consistent fashion with respect to their individual goals.

4.2. Pedagogical Goals

In addition to goals that represent the character profi le, our goal representation can encode the degree to which the pedagogy has been successful. We currently envision three approaches to encoding learning goals into Thespian.

First, learning goals can be embedded in the world's dynamics and the characters' goals. For example, one of the pedagogical goals in the MPE is for the student to learn to establish a relationship with the local people, in particular that they trust him/her. We can encode this pedagogical goal into the dynamics by ensuring that failure to establish trust will have consequences. At its most severe, distrust can cause irreparable breakdowns in the relationship. Specifi cally, in the MPE, if a student fails to achieve even the minimal requirement for this trust goal, the young man will accuse him of being a CIA agent, and all characters will refuse to talk to him. Such breakdowns are one extreme. Characters can also act in ways that help the student. In the MPE, the old man has the goal that he trust the learner, that he feel safe around him, and at times he deliberately behaves in a fashion that would elicit behavior from the student that increases trust. Specifi cally, the old man can ask the student questions about the student and his mission, which provides more information and makes the old man feel safer. Although it is not an explicit intention of the character, its behavior does assist the learner.

However, Thespian can provide characters with the explicit intention that the student learn. In this approach to encoding the pedagogy, characters have a goal that the learner acquire skills specifi ed by the pedagogy. A character could then use its mental model of the learner as a student model to measure the degree to which the pedagogical goals are achieved. The theory of mind embedded within Thespian forms a subjective view of the world that includes beliefs about the students' knowledge and capabilities based on their

behavior. The old man, for example, could have the explicit goal that the student give a high goal priority of establishing trust. Having encoded such a goal, the old man could now evaluate a possible action choice using its mental model of the student's goals to assess the effect on the student and, in turn, on the pedagogical goals so encoded. Again, because we have priorities on the goals, we can choose how much a particular character is driven by pedagogical goals for the learner in relation to its own personal goals.

Finally, a third way to encode the pedagogy is to have a behind-the-scenes director agent that is directing the drama in pedagogically appropriate ways. In other words, we could go even further by explicitly encoding the intention to teach in the overall system through this director agent. The MPE does not employ this technique currently but it is feasible within Thespian. These three approaches to encoding pedagogy (in the world's dynamics, in the character's intentions and in the system's intention) provide Thespian with a rich framework for realizing pedagogical drama.

4.3. Social Norms

While Thespian's ability to encode pedagogical goals gives the agents incentive to exercise the pedagogy, we also must give the *student* the same incentive. One of the motivations underlying IPDs is that the student's inherent social desires can provide an incentive for following the pedagogy if the characters are socially interesting entities. As described in Section 3, characters that are sensitive to social norms can provide such an incentive.

PsychSim provides a general framework for representing states, actions, and the dynamics of the world. While such probabilistic models have typically been used in modeling physical systems, Thespian uses them to model *social* dynamics as well. We constructed Thespian's model of social dynamics by first identifying critical social variables. We have begun by encoding the trust and liking relationships that exist between entities. We then defined domain-independent dynamics for these social variables (e.g., increase your liking of another agent if it does something that helps achieve your goals). Giving an agent goals on these social variables will give it incentive to be liked and trusted by the student. We are currently applying this same methodology to expand our set of social variables to include other key features (e.g., affinity, freedom).

In addition to these more persistent relationship variables, Thespian also uses social variables to represent more temporary obligations that may exist between agents. In general, actions by one agent can impose a type of obligation on another, and a certain set of responding actions will satisfy the obligation to some degree. We currently use these obligations to encode a broad set of social norms as pairs of initiating and responding actions: greeting and greeting back, introducing oneself and introducing oneself back, conveying information and acknowledging, inquiring and informing, thanking and saying you are welcome, offering and accepting/rejecting, requesting and accepting/rejecting, etc. For example, Thespian's dynamics for "inquiry" specify that one of its effects is the establishment of an obligation on the part of the inquiree to satisfy the enquirer (e.g., by providing the needed information).

By giving the agents goals to satisfy any such outstanding obligations, we give them an incentive to follow the encoded social norms. In some cases, the agents may already have an incentive from relationship goals in addition to the obligational ones. For example, an agent providing information in response to an inquiry will be helping the enquirer achieve its goals, leading to a stronger liking relationship. Alternatively, social norm goals may conflict with the agent's other goals, leading to possible violations. For

example, an agent may decide not to satisfy an inquiry obligation, because revealing the requested information may reveal vulnerabilities, threatening the security of the agent. The relative priorities among all of these goals reflect the value that the character places on the corresponding social norms. These values are often culturally specific and can also vary according to its personality. However, although we vary the relative weights on the norms from character to character, the underlying mechanism for representing and maintaining norms and obligations does not change, so we can reuse it across many IPDs.

4.4. Authoring

We have shown how Thespian encodes personalities, pedagogical goals, and social behaviors as goals that can drive autonomous agent behavior. Because of this autonomy, the author of the IPD no longer has to specify all of the possible behaviors of the character. However, the character's behavior now depends on the goal priorities chosen by the author, so we simply replaced the previous authoring task with a new one. Furthermore, the process of tuning such quantitative parameters is typically less natural to the author than writing a script.

Fortunately, PsychSim provides an algorithm for automatically choosing these goal priorities based on a few instances of desired behavior [11]. Thespian uses this algorithm to take *partial* scripts, provided by the author, and automatically tune the relative goal weights among the personal, pedagogical, and social goals of the character. Once Thespian has fit the character's goals to this input, the character will always generate autonomous behavior that is consistent with the given scripts, when applicable. Furthermore, when the learner's interactions lead them off the scripts, the agent will still act consistently with its goals. In other words, the fitting process extrapolates from the partial scripts to an exhaustive specification of consistent behavior over all situations. It is as if we were "teaching" the agent the motivations of its character, as opposed to having them simply memorize the scripts.

Thespian reduces authoring effort in two ways. First, Thespian's authoring process alleviates burden on authors by not requiring them to craft all possible paths through the story, while still allowing a more natural process than required by hand-tuning parameters. Second, Thespian supports the reuse of characters and environments across IPDs. Thespian can separate the models of characters from those of the environment they are in. Dynamics designed for one IPD environment can be reused in another. And after fitting, an agent becomes a character with a certain set of goals. This character can be easily plugged into other stories to play a similar role. See [12] for further discussion.

5. Results and Current Status

Figure 2 provides an excerpt of actual dialog between the (human) student, the student's aide and two locals from scene one. Note Figure 2 shows the surface language form, but this form is mapped by the speech recognizer to an internal speech act representation that the agents can reason about. Prior to this excerpt, the student has only introduced his name, but has not introduced the aide or details about their mission. The impact of having failed to make a proper introduction is that he has not built trust with the locals.

In the first line of Figure 2, the student asks the old man a sensitive question. However, the young man then seizes the dialog turn because he perceives a potential security threat by someone he does not trust. Through his mental model of the old man and his

Speaker	Addressee	Utterance
Student	Old man	*minu mas'uul b-hel-manTaqa?* Who is the most important offi cial in this town?
Young man	Old man	*9ala keefak!* Slow down!
Young man	Student	*u hiyye minu?* Who is *she*? (referring to aide)
Old man	Student	(silence)

Figure 2. Excerpt from Scene 1 dialog.

lookahead reasoning, the young man can foresee that if he does not stop the old man, the old man would give the answer to the student, which would hurt his own goal of safety. If he instead asks the student a question, he can not only stop the old man from giving the answer, but also gain safety by getting more information from the student. This reasoning leads the young man to tell the old man not to answer the question (second line from Figure 2) and to ask who the aide is (third line from Figure 2). The young man has both the goal of increasing safety and following social norms. According to the latter goal, he should keep quiet, because the student is asking the old man a question but he values safety more than following social norms. So, in this case, he picks the action that increases his safety, even if it violates social norms. For the old man, following social norms is the most important goal. He has two obligations. The student's question to him imposes an obligation to answer. The young man's question to the student imposes an obligation for the old man to keep quiet (i.e., wait for them to fi nish their conversation). The more recent obligation receives higher priority. Therefore, he chooses to keep silent.

This rather complex exchange was achieved by the automated fi tting process. Fitting adjusted the characters' goal weights (of safety and following social norms) to achieve the behaviors exhibited in this example.

Currently the MPE has three scenes. These scenes have as many as six characters plus the student's character. All three scenes are constructed by using automated fi tting. The TLTS system has so far undergone six stages of formative evaluation during the development process. We got mostly positive feedback about its effectiveness for language training. Since April 2004 to January 2005, we have gone through three rounds of testing with a total of 30 subjects. So far, the overall evaluation of the MPE is that it is successful in providing an engaging environment, and is an effective assessment tool [1]. Beginning in March of this year, we will have another round of testing with at least 100 subjects.

6. Conclusion

The promise of interactive pedagogical drama has often been thwarted by the arduous design and programming tasks facing the creators of such systems. Thespian facilitates the design process of agent-based IPDs in several ways. It enlists automation in the character confi guration process to simplify authoring. It also provides multiple ways to support pedagogical goals. Additionally, Thespian provides a methodology for modeling social dynamics within a decision-theoretic framework.

Thespian simplifi es the authoring process in several ways. Agents are motivated solely by their goals and their goals are automatically fi tted so that they perform according to the scripts. Because their behavior is driven by their goals and not simply scripted, the agents respond to unexpected user interaction in ways consistent with their motivations. If they do not, the misbehavior can also be fed into the fi tting process. In

the MPE, we have demonstrated how to embed pedagogical goals in the dynamics of the story world and have discussed additional approaches. We believe these techniques can be applied to other IPDs as well.

Going forward, the vision of Thespian would be for non-technical designers to author dramas on their own. There are still steps in the process that are impediments to such a vision, including translating scripted dialog into the formal speech act language that the agents understand. We plan on addressing such impediments in our future work.

Acknowledgments

This project is part of the DARWARS Training Superiority Program of the Defense Advanced Research Projects Agency. The authors wish to acknowledge the contributions of the members of the Tactical Language Team.

References

[1] C. Beal, W.L. Johnson, R. Dabrowski, and S. Wu. Individualized feedback and simulation-based practice in the tactical language training system: An experimental evaluation. In *AIED*, 2005.

[2] A. Chulef, S.J. Read, and D.A. Walsh. A hierarchical taxonomy of human goals. *Motivation and Emotion*, 25:191–232, 2001.

[3] W. Swartout et al. Toward the holodeck: Integrating graphics, sound, character and story. In *Agents*, pages 409–416, 2001.

[4] W.L. Johnson et al. Tactical Language Training System: An interim report. In *Proc. of the Internat'l Conf. on Intelligent Tutoring Sys.*, pages 336–345, 2004.

[5] I. Machado, A. Paiva, and P. Brna. Real characters in virtual stories: Promoting interactive story-creation activities. In *Proc. of the Internat'l Conf. on Virtual Storytelling*, 2001.

[6] S. Marsella, W.L. Johnson, and C. LaBore. Interactive pedagogical drama for health interventions. In *Proc. of the Internat'l Conf. on Artificial Intelligence in Education*, 2003.

[7] S. Marsella, D.V. Pynadath, and S.J. Read. PsychSim: Agent-based modeling of social interactions and influence. In *Proc. of the Internat'l Conf. on Cognitive Modeling*, pages 243–248, 2004.

[8] M. Mateas and A. Stern. Integrating plot, character and natural language processing in the interactive drama Fa cade. In *Proc. of the Internat'l Conf. on Tech. for Interactive Digital Storytelling and Entertainment*, 2003.

[9] A. Paiva, J. Dias, D. Sobral, R. Aylett, P. Sobreperez, S. Woods, C. Zoll, and L. Hall. Caring for agents and agents that care: Building empathic relations with synthetic agents. In *Proc. of the Internat'l Conf. on Autonomous Agents and Multiagent Systems*, pages 194–201, New York, 2004. ACM Press.

[10] L. Plowman, R. Luckin, Laurillard, M. Stratfold, and J. Taylor. Designing multimedia for learning: Narrative guidance and narrative construction. In *CHI*, pages 310–317, 1999.

[11] D.V. Pynadath and S. Marsella. Fitting and compilation of multiagent models through piecewise linear functions. In *AAMAS*, pages 1197–1204, 2004.

[12] M. Si, S. Marsella, and D.V. Pynadath. Thespian: Using multi-agent fi tting to craft interactive drama. In *AAMAS*, 2005.

[13] R.D. Smallwood and E.J. Sondik. The optimal control of partially observable Markov processes over a fi nite horizon. *Operations Research*, 21:1071–1088, 1973.