Experiences authoring interactive pedagogical dramas

Stacy C. Marsella

Information Sciences Institute,
University of Southern California,
4676 Admiralty Way,
Suite 1001,
Marina del Rey,
CA 90292, USA
E-mail: marsella@isi.edu

Abstract: The focus of this article is the design of Interactive Pedagogical Dramas (IPD). An IPD is a computer-generated environment that immerses the learner as an active participant in an engaging, evocative story, populated with animated characters. The story unfolds based on the learner's decisions. The various design issues that are faced in crafting IPDs will be covered. In particular, how pedagogy can be incorporated into the learner's interaction with the narrative will be discussed. The discussion will be illustrated using several existing IPDs. Finally, we will conclude with several observations on the design process for IPDs and possible new directions for this design process.

Keywords: generative learning; interactive narrative; interactive pedagogical drama.

Reference to this paper should be made as follows: Marsella, S.C. (2008) 'Experiences authoring interactive pedagogical dramas', *Int. J. Continuing Engineering Education and Life-Long Learning*, Vol. 18, No. 2, pp.159–180.

Biographical notes: Stacy Marsella, PhD, is a Project Leader at the Information Sciences Institute and a Research Assistant Professor in Computer Science at the University of Southern California. His research interests include interactive drama, the computational modelling of social interaction and the application of technology to education.

1 Introduction

Drama is a powerful pedagogical tool. In poetics, Aristotle argued over two millennia ago that learning and drama are interwoven: that drama is an imitation of life and humans learn through enjoyment of that imitation. More recently, research in narrative psychology has argued that narrative is central to how we understand the world and communicate that understanding (Bruner, 1990), and of course, the engrossing motivational nature of story is unmistakable; the world now consumes stories in various media with a 'ravenous hunger' (McKee, 1997).

Yet stories can also have a drawback from a learning perspective: they typically place the learner in the role of passive audience instead of active learner. In contrast, the goal of

Interactive Pedagogical Drama (IPD) is to exploit the edifying power of story while promoting active learning. An IPD is a computer-generated environment that immerses the learner as an active participant in an engaging, evocative story. The story environment is realised as an animated world, populated with virtual characters. The learner makes decisions or takes actions on behalf of a character in the story or can play a more directorial role in deciding the direction of the overall story. Regardless, the learner experiences the consequences of her decisions. Ideally, the learner identifies with and assumes responsibility for the characters in the story, while the control afforded to the learner enhances intrinsic motivation (Lepper and Henderlong, 2000), and since the IPD framework allows for stories with multiple socially interacting characters, learning can be embedded in a social context (Vygotsky, 1978).

IPDs have been developed to address a range of educational challenges. Carmen's Bright IDEAS (CBI) (Marsella, Johnson and LaBore, 2000; Marsella, Johnson and LaBore, 2003; Marsella, Gratch and Rickel, 2004) is an interactive, animated health intervention designed to improve the social problem solving skills of mothers of paediatric cancer patients. The learner influences the behaviour of the main character, a virtual mother who is facing similar stresses to the learner. The Tactical Language Training System (TLTS; Johnson et al., 2004) is designed to teach the basics of a foreign language and cultural awareness. The learner navigates in the virtual world, where she is the main protagonist and interacts with the virtual characters using spoken language and gestures. FearNot! (Paiva et al., 2005) is designed to teach children about adolescent violence and how to deal with it. In the virtual environment, there is a bully, Luke, who is tormenting John. The learner provides advice to John who then may act on that advice.

IPDs are, in many ways, a high-tech approach to the role playing games used, for example, in child education (Shaftel and Shaftel, 1982), therapy (Yardley-Matwiejczuk, 1997), conflict resolution (Rasmussen and Oatley, 1992) as well as for entertainment (e.g. Dungeons and Dragons). However, IPD alleviates the need for multiple human participants in a physical performance space by creating a virtual world populated with virtual embodied characters. The use of a computer-generated environment also allows more precise control over the learner's experience, supports embedded assessment of their progress and can provide a rich sensory environment that is both engaging and memorable. IPDs are also closely related to the work of Schank (1996), Klein (1998), and others on the edifying narratives of personal experience (war stories). However, such narratives are fundamentally less interactive. Some story-oriented learning methods (e.g. Jasper series (Cognition and Technology Group at Vanderbilt, 1992)) promote active discussing and problem solving in the context of stories, but do not provide learners a way of seeing the consequences of their decisions.

The design of interactive pedagogical drama faces several challenges. The effort required to design and build interactive dramas can be quite significant, 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.

The most fundamental challenge for an interactive pedagogical drama, however, is that it must achieve its pedagogical goals. This leads to the question of how the narrative of an IPD relates to the pedagogy and achieves desired learning outcomes. Narrative has been defined as the "semiotic representation of events meaningfully connected in a temporal and causal way" (Onega, 1996). We take a strong view of how these narrative connections between events relate to the pedagogy. We presume that an IPD should strive to causally connect events in the drama, including character behaviour, in ways that are consistent with and reveal the pedagogy. For example, if the goal is to teach how stress influences problem-solving, as in CBI, then there will be connections in the narrative between a character's stress, the decisions made and the ramifications of those decisions. Further, this relation between narrative events and pedagogy should be maintained regardless of how the story unfolds based on the learner's actions. This is particularly required for the key moments of high drama.

We call this tight connection, between narrative and pedagogy, a *strong narrative entanglement*. Although it has typically guided the author's work, there are weaker ways to entangle narrative and pedagogy. For example, the narrative can simply be used to entice or motivate a learner to interact with the system while learning is based on activities that are not part of the narrative. One consequence of strong narrative entanglement is that the events and the characters become simulations of the pedagogy and theories that inform the pedagogy. We will discuss in this article how this can have a significant impact on the design and technology of these systems.

In addition, the design must also address how the learner is entangled with the narrative. The designer must consider the role of the learner in IPD and whether they are even present as a character in the narrative. The designers of CBI, TLTS and FearNot!, for example, each gave the learner a very different role in IPD. In this article, we will seek to explicate these design choices by arguing for what we call *pedagogically mediated involvement*, specifically that the learner's role, sense of presence in the environment and their affective response to it need to be mediated by pedagogical concerns. As part of pedagogically mediated involvement, we will consider not only the role given the learner but also the factors that mediate the presence such as the learner's sense of control over the environment, the visual setting, the sounds, etc. These aspects all become critical design decisions in crafting the learning experience. This view that IPDs should seek to mediate and potentially reduce the learner's sense of presence distinguishes IPDs from virtual environments for games, films and other related forms of media used for entertainment where achieving high presence is often a central goal.

In addition to the numerous narrative design choices, there is also the choice of technology that realises the interactivity. A range of technologies have been employed to create interactive narratives. For example, Read et al. (2005) used a *branching narrative*, in their interactive video to reduce risky sexual behaviour. The learner decides between one of a few alternative choices at fixed points in the story and the plot branches based on that decision. The action that unfolds between branches is typically a canned presentation such as a pre-recorded video or animation segment. The main drawback of branching narratives is that it is very difficult to provide a high degree of interactivity. More interactivity requires more branches. With more branches, it rapidly becomes too burdensome to manage all the paths through the branches, ensure the story is coherent regardless of path and craft all the dialog, video or animations required. In practice, this has limited the amount of interactivity.

As a consequence, researchers in interactive narrative have been exploring a range of generative approaches that do not require the designer to explicitly craft all possible paths through the interactive drama. One common approach in the design of IPDs is to cast interactive drama from the perspective of the characters in the drama. Along the lines of role-playing games, this approach uses a multi-agent framework where autonomous software programmes called agents play the role of the main characters. The agents (often called Virtual Humans or Embodied Conversational Agents (Cassell et al., 2000)) are self-motivated, with their own goals. The narrative follows from the characters seeking to achieve their goals and how those characters may cooperate and compete to achieve those goals. A key benefit of this approach lies largely in the fact that goal-motivated agents can provide coherent character behaviour even in the face of responding to open-ended user interactions. The key challenge for this approach is to constrain the character behaviour so that it is also consistent with pedagogical goals.

Finally, there is an issue of the design process. IPDs are complex artefacts with a design process that must coordinate the pedagogy, narrative design, software technology and artwork. The author has been involved in the creation of six IPDs. Although the systems were very different, their development shared common elements in how the challenges of the design process were approached. In particular, the process can be characterised as having the following steps. The pedagogy establishes what is going to be taught. A scriptwriter then becomes immersed in the pedagogy or works with pedagogical experts directly and then crafts a story concept and script. Although the resulting stories have typically been linear, non-interactive narrative story, they serve as a solid starting basis for creating an interactive experience that ties together dramatic and pedagogical goals. The stories have had a largely traditional story structure, including some opening scenes or back-story that establish the characters, scenes that establish the conflicts or dilemmas faced by the protagonist and then some resolution of the conflict or dilemma into a positive or negative ending. The story is then realised in the interactive design and technology while an art team concurrently works on the animations for the character bodies and the background scenes.

The focus of this article is the design of IPDs. The various design issues that are faced in crafting IPDs will be discussed in detail, from a perspective admittedly greatly influenced by personal experience. In particular, the impact that strong narrative entanglement and mediated learner involvement have on the technology will be discussed in greater detail. Then three systems will be covered. In addition to CBI and TLTS mentioned above, Stability and Support Operations (SASO; Traum et al., 2005) will be discussed, a research prototype where the learner engages in a spoken language negotiation with a life-sized virtual human. Finally, we will conclude with several observations on the design process.

2 Design

The central issues for the design of an IPD are the nature of the drama, the relation of the drama to the pedagogy and the role that learner plays in that drama. In this section, these issues will be considered in greater detail.

2.1 The drama

There is a rich history of traditional literature from which we can draw inspiration on crafting interactive narratives. Taking a genre-based view, we might consider action, adventure, mystery, fantasy or psychological drama as the basis for the pedagogical experience. Often this decision is constrained by concerns over transfer of the learned skills. Specifically, TLTS, Carmen and FearNot! all use stories that closely mirror the real world problems that the learner will be facing, in part due to the presumption that this will facilitate transfer. There is also a wide range of narrative, cinematic and theatrical techniques that can be considered. For example, will the experience adhere to real time so events unfold at the same pace as they do in the real world (e.g. as in the film High Noon)? Will the narrative have a linear structure that proceeds forward in time or will it be more unhinged from time, wandering forwards and backwards (e.g. Vonnegut's Slaughterhouse Five)? Various non-linear techniques can be blended into a largely linear structure, such as flashbacks to reveal prior events and flash-forwards to envision alternative possible outcomes.

2.2 Narrative entanglement

Another key design choice is the degree to which the narrative is entangled with the pedagogy. At one extreme, the narrative of an IPD may be used only as a means to facilitate engagement, with the pedagogy grafted on largely as interlude. Along these lines, one might imagine an adventure story, a group of children on a space mission in which problem solving interludes are interjected that teach basic physical laws about mass and momentum. The problem that the learner faces may be quite disconnected from the plot content of the story but mastery of the skill or knowledge is required before the learner is allowed to return to the adventure. Somewhat more entangled, the problem may arise in the context of the plot. In contrast, in the case of strong entanglement, the basic events of the story may be driven by that knowledge. For example, movement of the spaceship may be a simulation of those physical laws and therefore whenever the learner tries to alter the navigation of the ship, they face a potential learning experience. Such tight merging of story and simulation of the pedagogy is not limited to the physical world. One can, for example, craft a story that teaches around a social simulation that adheres to a theory of social interaction. Of course, this issue of pedagogical entanglement tends to be coupled to the issue of whether the setting of the drama is a facsimile of a real world setting in which learned behaviour is to be realised.

In the author's work, strong entanglement has largely been viewed as an ideal for IPDs. IPDs are designed to be experiences consistent with constructivist theories of learning (Bransford and Vye, 1989). Specifically, we seek to craft a narrative-based environment in which learning outcomes are achieved through active involvement in the story and the experiences that come out of that involvement. The experiences should reveal and dramatise the pedagogy. In our work, this has meant that the causal and temporal connections between narrative events are largely rooted in the pedagogy. Narrative events happen the way they happen due to the pedagogy. Further, moments of high dramatic conflict or resolution in the narrative have the potential for making the revelation more vivid and memorable for the learner. Dramatic moments that are unrelated to the pedagogy are, in this view, missed learning opportunities that may even obscure the pedagogy. This tight entanglement of narrative and pedagogy has several ramifications. Clearly, it may constrain how the drama unfolds. Moreover, the underlying technology that determines how events unfold must embody this constraint. It must in some way computationally model the pedagogy so that the learner can experience pedagogically meaningful connections between narrative events regardless of how the story unfolds. For the agent-based approaches to IPD where the story unfolds based on how the autonomous characters react, the characters must therefore embody, behave consistent with, the pedagogy.

2.3 Pedagogically mediated involvement

The learner can play a variety of roles in IPD. They can be the main protagonist with direct responsibility on how the story unfolds. In particular, in SASO and TLTS, the learner is the main protagonist and the goal is to place the learner under stresses in the virtual world similar to those they would have in the real world. At the same time, this ceding of control to the learner can make overall control of the story's direction more difficult. Alternatively, the learner can also be a supporting character. The learner in such a role may have much less control over, and therefore less responsibility for, how the story unfolds. This places some psychological/emotional distance between the story and the learner. In the FearNot! Story, for example, the learner is a trusted advisor to the main character, a child victimised by a bully. The learner is never part of the main action scenes. Rather the main character comes to the learner and asks for advice from the learner. The advice then potentially influences the main character's behaviour in subsequent scenes. Similarly, the learner may be entirely off-screen, exerting a more backstage indirect directorial influence on the characters and events.

IPDs are multimedia experiences with animated characters. Thus a key issue becomes the design of the background artwork, the characters and the animation. These media choices play an important role in the learner's sense of presence in the environment as well as their affective responses to it. There is a large space of design variation here. The SASO system sought to immerse the learner in a very realistic looking, even stress-inducing experience. There is a gritty-ness to the background designed to convey a war-torn world. In contrast, CBI has two-dimensional (2D) cartoon-like characters, embedded in a highly stylised environment designed to be calming. As we will discuss, these differences in learner's role, interactions with the environment and presence in the environment were deliberate design choices determined in large part on what was being taught and decisions on how it should be taught.

2.4 Approaches to realisation

The discussion in this article will focus on an autonomous agent-based approach to interactive pedagogical drama, where autonomous agents, 'self' motivated software programmes, realise the characters in the story. These agent-based approaches focus more on the characters and how to constrain their behaviour so that it is consistent with narrative. However, designers have come at the challenges of making interactive narratives from a variety of perspectives (Weyhrauch, 1993; Galyean, 1995; Cavazza, Charles and Mead, 2001; Gebhard et al., 2003; Riedl, Saretto and Young, 2003; Szilas, 2003; Machado et al., 2006). In particular, one can focus more on the narrative structure instead of the characters and use a computational model of narrative form to adapt the user experience in the face of interaction. For example, Mimesis (Young, 2001) exploits a plan based representation of narrative that supports revisions based on user interaction. Façade (Mateas and Stern, 2003) breaks down the narrative structure to dramatic beats that can be strung together and modified by user actions. Façade and Mimesis take a view of the interactive narrative problem whereby events must obey some overarching structure and it is adherence to that structure that ideally will allow the player to form meaningful connections between events. We see this as a centralised, largely top-down approach to the interactive narrative problem.

We take a different approach to the interactive narrative problem. Borrowing ideas from improvisational theatre, role play, multi-player games, ensemble cast movies and simulation-based training, we recast interactive drama from the perspective of the characters in the drama. In this more decentralised, character-based view of interactive narrative, the player's ability to coherently interpret the narrative events follows from the goal-directed behaviour of the characters seeking to achieve their goals and how those characters may cooperate and compete to achieve those goals.

Such a distributed, more autonomous character driven approach essentially flips the design standpoint from adding variability into a model of the story to constraining variability of a multi-agent system. Ideally, such a change in design stance promises designs that support more open-ended interactions, increase player's sense of control and potentially increase engagement.

The operation of our agent-based interactive pedagogical dramas involves several components: a cast of autonomous character agents, the 2D or 3D puppets which are the physical manifestations of those agents, potentially director and/or cinematographer agents, and most importantly a role for the learner/user to play in the interactive drama. Animated agents in the drama not only choose their actions autonomously but also follow directions from the learner and/or a director agent. Director and cinematographer agents can be used to manage the interactive drama's onscreen action and its presentation, respectively, so as to maintain story structure, achieve pedagogical goals, and present the dynamic story so as to achieve best dramatic effect. The design of all these agents requires both general capabilities as well as knowledge specific to the interactive drama that is being created. We have typically approached the problem of developing these components by deconstructing the story line into the specifications for these agents.

3 Carmen's Bright IDEAS

We begin the discussion with the earliest system the author was involved in designing, CBI. CBI is an interactive, animated psycho-social health intervention. The system was designed to improve the social problem solving skills of mothers of paediatric cancer patients. These mothers face numerous stressful problems in addition to their child's illness. There may be problems at work due to having to take time off to bring their sick child to the hospital. Their other children may feel neglected and scared because their sibling is seriously ill. CBI taught these mothers a methodical, step-wise problem solving technique called Bright IDEAS to help them address such problems (Varni et al., 1999). IDEAS is an acronym for Identify the problem, Develop solutions, Evaluate the solutions, Act on the solution and See if it worked. Bright stood for the importance of having a positive outlook and believing in one's own self-efficacy. Through the process of interacting with the story by selecting thoughts for its main character, the mother was introduced to Bright IDEAS and experienced its application to Carmen's problem, a virtual character that had problems that mirrored the problems of the real world mother (Figure 1).

CBI is a 3-act drama. The first act provides the back-story of what difficulties the main protagonist, Carmen, is facing due to the illness of her son, Jimmy. This back story includes the problems that Carmen is having with her healthy daughter Diana who is feeling scared and neglected, with her boss because she is taking a lot of time off to take Jimmy to the hospital and with understanding the doctors. The second act concerns

Carmen interacting with Gina, a clinician. Gina is going to teach Carmen a problem solving technique for dealing with her stressful, social problems, called Bright IDEAS. Gina guides Carmen through the application of Bright IDEAS to one of her problems (which one is selected by the user). The third act reveals what happens when Carmen applies Bright Ideas to some of her problems.

Figure 1 Scenes from Carmen's Bright IDEAS

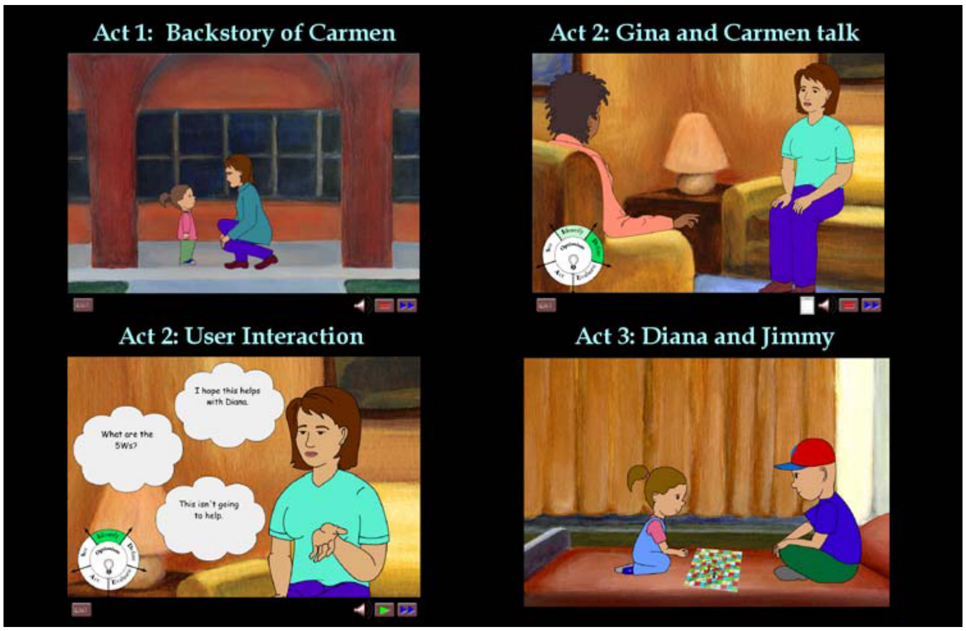

The second act is the central part of the interactive experience. The user influences Carmen by selecting thought balloons which in turn impact Carmen's internal emotion and cognitive models and therefore her subsequent behaviour. Meanwhile Gina can also influence Carmen emotionally and cognitively via her dialog. Thus both the user and Gina influence Carmen. While the user is free to pull Carmen's coping in helpful or harmful ways, even possibly getting her to act out, Gina is attempting to use dialog to get Carmen back on track to a successful pedagogical and dramatic outcome.

3.1 Pedagogy and narrative entanglement

The pedagogy of the Bright IDEAS technique was informed by Richard Lazarus's appraisal theory (1991), a theory on the causes of emotion and stress as well as how people cope with emotional stress. In Lazarus's theory, a person's emotions arise from a subjective appraisal of situations along several dimensions such as whether it is congruent with the person's concerns, whether they have sufficient control to improve the situation, etc. This appraisal in turn determines a coping response that attempts to regulate the emotional response. If a person appraises that they have control, for example, they will tend to cope by planning, or taking action, to improve a bad situation. If they appraise that they don't have control they will use other coping responses such as try to avoid thinking about a bad situation, engage in wishful thinking that problems will go away, become resigned that the situation will just have to run its course, etc. As the situation evolves and as a person copes with it, there will be subsequent re-appraisals and coping responses.

The Bright IDEAS problem solving technique seeks to alter the learner's appraisals and coping responses to situations. The IDEAS step-wise problem solving technique ideally gives the learner a systematic way to appraise their control over a situation, allowing the person to better judge the problems that can be solved from those that cannot and to increase their sense of self-efficacy in solving problems that are in fact solvable. Similarly, the techniques stressed maintaining a positive (Bright) attitude about one's own self-efficacy in solving problems. In addition, the pedagogy stressed that sometimes initially distancing oneself or letting go when problems are too difficult can actually lead eventually to more effective problem solving, as part of the re-appraisal process.

The interactive narrative of CBI directly reifies this pedagogy in the agent models. Specifically, Carmen is largely driven by a computational model of Lazarus's appraisal theory and Gina instantiated Bright IDEAS. Separately, Gina's and Carmen's behaviours were designed to vividly illustrate Bright IDEAS and appraisal theory, respectively. Their interactions are designed to deeply model and ideally reveal how Bright IDEAS works on a person's cognitive emotional processes.

Specifically, the agent architecture used in CBI has both a dialog model and an emotion model. The emotion model was designed around a computational model of Lazarus' appraisal theory. The emotion model in turn informed the agent's dialog choices. In Carmen, these choices were largely a function of her emotional state and how she was choosing to react to that state. In the case of Gina, the dialog model represented dialog as a hierarchically structured goal-oriented interaction. Certain tasks had to be performed in the interaction and there were alternative strategies (recipes) for achieving these tasks. The task knowledge embedded in Gina's dialog model included the steps of Bright Ideas and as well as a range of abstract strategies for persuading a person to go through a joint problem solving effort (see Marsella, Gratch and Rickel (2004) for a more thorough discussion of the agent architecture).

3.2 Pedagogically mediated involvement.

A design goal for the interactive experience was to provide a safe environment that facilitates exploration. These mothers are clearly under enormous stress in general and would actually being using the application when they took their children to the hospital. Such stress can aversely impact attention, problem solving and learning. Often a key aspect of dealing with stressful encounters is to achieve some distance from the problem to reduce stress and facilitate subsequent problem solving, a point that the pedagogy highlighted. Thus, a design decision was made to manage the stress that would arise from using the application. We wanted a relative low stress application, a respite that would facilitate the freedom to explore coping strategies within the application.

In keeping with the stressful situation the learner is undergoing, and the pedagogical goal of providing some distance for the learner, the interaction was designed as a presentational, third person experience, not an immersive experience. The learner did not play one of the characters in the story or control them moment-to-moment since they were self-motivated. The characters also did not directly talk to the learner. The learner only influenced Carmen through selecting thought balloons (Figure 1) and was therefore not directly responsible for what happened. The discussion was about Carmen's problems, not the learners. Again this was a deliberate distancing as well as avoiding the possibly insulting situation of a cartoon character talking about a real mom's serious problems. Interactivity was kept low, again deliberately to reduce stress.

In addition, CBI had a cinematography agent that controlled the camera. The agent had a collection of simple rules designed to enhance engagement in the experience, promote pedagogy and ensure Carmen's emotional state was easily decoded by the learner. For example, these rules were designed to do close-ups when the agents' emotion models revealed strong emotions. To highlight pedagogical points and make for a more engaging experience, manipulation of time and time compression were used liberally. The first act cuts between several past episodes that provide back-story on Carmen's problems. In the second act, flashbacks are used. More interestingly, hypothetical, what-if 'flash-forwards' are used that allow the learner to see possible future outcomes of Carmen's decisions.

The 'realism' of the underlying agent models is also tweaked for the experience. Carmen's emotional model, although based on theories of human emotion, operates with an unrealistic 'low-mass' dynamics in order to reveal her inner conflicts. For example, Carmen can dismiss Diana's tantrums as being just 'childish' and then feeling guilty about dismissing them all in the space of 1–4 dialog turns, depending on Carmen's emotional state and whether Gina needs to prompt her for further explanation. Such a rapid transformation is likely not realistic but serves to provide both an engaging and revealing experience. Similarly, Gina's strong impact on Carmen's emotional state, by re-assuring her for example, is also somewhat unrealistic but useful for ensuring certain positive outcomes could be forced in the drama.

Finally, the art design was also strongly informed by the pedagogy and learning context. The characters were 2D – slightly divorced from reality. The backgrounds appeared as if they were shot through a haze. The learner saw a calming Edward Hopper-esque living room like environment with muted colours and lighting that is quite different from the actual hospital settings that the learners faced when talking to clinicians in the real world. All these artistic elements were geared towards calming the learner and providing psychological distance from the experience.

3.3 Design process

The above design elements of the system were informed by several formative evaluations with target learners. The facial expressions were evaluated in terms of how readily people could interpret the emotion being expressed. The 2D bodies also served to greatly simplify decoding of expressions by admitting simpler expressions and gestures. Carmen herself was designed to promote identification with the target learners. Her ethnicity and age was supposed to be somewhat ambiguous but nevertheless overlap the target audience.

The well-defined pedagogy gave an excellent basis to write the script that tightly integrated the pedagogy. The story itself was designed to mirror the kinds of situations that the learners were facing in the real world. The nature of the problems that our target learners faced was well-known and further refined in exploratory design interactions with representative target learners. This essentially defined the basic suite of social problems our virtual mom, Carmen, would face. Act 1 revealed those problems. The application of the Bright IDEAS technique formed the background of the Act 2 interaction between Gina and Carmen, and Act 3 revealed the resolution of Carmen's problem solving attempts. Further, the scriptwriter was adept at writing dialogue and creating characters that were well and consistently motivated.

3.4 CBI discussion

Overall, the CBI experience was designed to mirror the kinds of problems the learner faces but be distinct from those problems. Rather the learner was supposed to be transported into the drama away from their own problems. Whereas the emotions evoked would be strong, the transportation would make them more manageable, important considerations given the stress the learner is under.

The story and interactive technology took into account a wide range of factors such as the underlying psychological theories that the Bright IDEAS technique was based on, the problems faced by the target learners, their stressful situation, etc. The design also took into account that English may be a second language and computer experience would be low, as was the case for the two clinical sites involved in designing Carmen.

The key conclusion to be drawn here was the central importance of the pedagogy in all phases of system design. It allowed for a largely top-down design that gave critical guidance to story development, the design of the interactive technology and the artwork. Nevertheless, the process of designing the system was highly technical. As a consequence, the story's writer could not play a direct role in designing the system, though he did influence it indirectly. The story's crafting of well-motivated characters, in particular, simplified the subsequent agent design.

4 SASO

Figure 2 depicts SASO. SASO is part of an overall research agenda of creating a general virtual human architecture that can engage in spoken language interaction with humans. The project has roots in the STEVE system (Rickel and Johnson, 2000) and the Mission Rehearsal Exercise (MRE) (Rickel et al., 2002; Swartout et al., 2005). The basic application envisioned for these virtual humans is the creation of social training environments where a student could learn to deal with high stress social interactions in the safety of a virtual world. SASO is a scenario in which the trainee must negotiate with a doctor running a clinic. A life-sized virtual human plays the role of the doctor. The human trainee negotiates with the doctor to get him to move the clinic, which could be damaged by a planned military operation. Ideally, the trainee will convince the doctor. Figure 2 shows the trainee's view of the doctor in his office inside the clinic. The success of the negotiation will depend on the trainee's ability to follow good negotiating techniques, when confronted with different types of behaviour from the virtual doctor.

Figure 2 The doctor in SASO

4.1 Pedagogy and narrative entanglement

One of the central ways to characterise negotiation under adversarial conditions is with respect to the tension between competition and cooperation. Negotiators may have different goals, perceive themselves in conflict over those goals but may also perceive the need to cooperate to some degree to achieve their goals. In this view, one can characterise the state of a negotiation process from the perspective of this competitive/cooperative orientation of the negotiating parties to the negotiation and the strategies they employ in light of those orientations. Specifically, one oft-made distinction is integrative vs. distributive (Walton and Mckersie, 1965; Putnam, 1990) orientations. If negotiators perceive a negotiation as a win–lose game where there is a fixed value to be distributed, then it is as called distributive. There will be a winner and a loser. In contrast, an integrative situation is one where both sides can potentially win, a win–win situation where negotiation could add value and be of benefit to both sides. These basic distinctions presume some commitment to engage in negotiation. However, an individual may simply have an orientation that there is no possible benefit or even need to negotiate. Their approach to the negotiation may simply be to avoid it or deny the need for it, what is termed avoidance (Sillars et al., 1982).

People tend to use a range of dialog tactics consistent with their orientations (Putnam and Jones, 1982). Avoidance tactics include shifting the focus of the conversation and delaying tactics. Distributive tactics can include various defensive moves such as stating prior commitments that bind the negotiator or arguments that support the negotiator's position. Distributive tactics can also be more offensive, such as threats, criticisms, insults, etc. Integrative tactics are more cooperative with negotiators actually attempting to see issues from the other's perspective. Tactics include arguments that support the other's position, acceptances of offers, offers of support, etc. Note at a finer grain of analysis, the tactics employed have both instrumental and affective components. For example, distributive tactics, besides trying to gain competitive advantage, tend to convey or realise a negative affect situation whereas the integrative tactics try to promote a positive affective climate.

Negotiation can move in and out of these orientations, though the ideal of course is for the negotiators to move towards integrative tactics. Several factors have been identified as being critical to moving towards integrative tactics, including acts of reciprocity, establishing trust, reinforcing shared goals, etc. (Wilson and Putnam, 1990).

These orientations and associated dialog tactics formed the basis of SASO's doctor. Unlike CBI, the virtual human architecture at the heart of SASO already existed as part of the earlier MRE project. The centrepiece of the virtual human architecture is a task model that gives the virtual human a declarative representation with which to reason about its goals, how to achieve those goals and the role of other people in helping or hindering the achievement of goals. The architecture also includes a model of emotional reasoning, **Em**otion and **A**daptation (EMA) (Gratch and Marsella, 2004) that was based on the same cognitive appraisal theories that influenced CBI's agent models but is a more complete model of appraisal and coping. There is also an advanced dialogue reasoning component (Traum and Rickel, 2002). The dialogue, emotion and task reasoning are closely coupled, allowing specifically appraisal and coping processes to be influenced by the interpretation of dialog and to influence the virtual human's action and dialog choices.

In the case of the doctor, the task model represents that moving the clinic has ramifications in terms of his goals and tasks. Will there be supplies at the new location, will they be able to move the very sick patients and will ambulatory patients be willing to travel to the new location? Because of the dialogue component, he can talk about those ramifications and due to the emotion model he can emotionally react to them.

The distinctions identified in the above research on negotiation had rather direct correspondents to our virtual human architecture. The negotiation orientation can be recast into appraisal theory and is therefore consistent with the appraisal/coping process in the architecture. Thus, the negotiator's perception of the negotiation situation is in principle an appraisal of the situation and their orientation or strategy towards the negotiation is in principle a coping strategy. Specifically, an avoidance negotiation orientation arises from an appraisal that the negotiation is undesirable but avoidable. That leads to an avoidance coping strategy. A distributive orientation comes about if the situation is now deemed unavoidable. That leads to coping strategies of attacking or defending against the threat. Integrative orientation arises due to appraisals that suggest potential benefits and this in turn leads to more constructive problem solving. Movements between negotiation orientations thus become a part of the architectures appraisal, coping and re-appraisal processes. The dialog management and natural language generation use the coping strategy recommendations to inform its selection of dialog tactics, such as attempts at shifting topics.

4.2 Pedagogically mediated involvement

The learner is immersed in the experience as the main protagonist. In fact, the goal of SASO, in stark contrast to CBI, is to reduce the psychological distance between trainee and experience, to put the trainee into a realistic setting and to place them under similar stresses they would face in similar situations in the real world. The overall experience is designed to give a sense of immersion in linear time where events are supposed to unfold as they largely do in a real world experience. There are, for example, none of the flashbacks or flash-forwards of CBI.

As can be seen in Figure 2, SASO's art design is based in realism, again in stark contrast to CBI. SASO also has injured people in the clinic, a woman crying over the loss of her husband, etc. The learner is immersed in this realism. One key assumption of this experience is that it will evoke the emotions and stresses of the real crisis negotiation and evoking those emotions makes a more effective environment for teaching negotiation. The environment seeks to teach emotion regulation because of the importance it plays in a real world emotion-charged negotiation, where it is important to regulate emotional reactions, deal with the stresses, and not lose control in the face of one's own emotions and the emotions of the other party in the negotiation. Further, there is also an assumption that realism will facilitate transfer of learned skills to the real world. These assumptions are common in virtual reality training systems. There often is a tendency to err on the side of realism in such applications even where there is uncertainty on how important realism may be. Nevertheless, the contrast with the approach taken in CBI is striking.

4.3 Design process

The project had as a goal at the onset to engage the writers of the script in creating the variability of the interactive experience by educating them in the virtual human architecture the project was using. Several writers and a game designer worked on the project and there was an active effort to engage them in designing interactive variability of the experience by informing them of the virtual human's capabilities. However, the final script idea actually came from the agent designers.

The advanced architecture greatly facilitated adapting the negotiation pedagogy to the virtual human's cognitive-emotional models. In fact once a script was in hand, the implementation of the scenario took little time, roughly on the order of weeks to encode the script into the virtual human's task model. On the other hand, it was quite difficult for the writers to fully grasp the capabilities of the architecture and exploit it in designing interactive variability. As a result, there were delays in getting the script. The script that was eventually written was in fact based on an idea of one of the agent designers, working in reverse from the way it worked in CBI. At the end of the writing process, much like in CBI, the agent designers still only had largely linear scripts to inform agent design. These scripts played only a small role in agent design, restricted largely to informing some of the dialog used.

4.4 SASO discussion

Several conclusions can be drawn from the SASO project. The virtual human architecture is powerful and flexible. In the hands of knowledgeable designers, it was straightforward to model a new scenario. Once the script was written, it took only 90 days to craft the entire SASO scenario (agents and artwork). However, this powerful architecture also made it difficult for writers to fully grasp the capabilities of this new agent-based interactive media and they could not effectively create for it. As a result, the flow in this design process was largely the reverse of CBI, from the technology design to the writing. Even the pedagogy was based on the agent designer's research and was interpreted in terms of the existing virtual human architecture, though in this case it is more a testimony to the flexibility of the architecture. Like CBI, authors were still not involved in designing interactivity.

Subsequent to SASO, the authoring of these virtual humans is now being addressed by creating authoring tools that facilitate the job of non-experts creating task models for the virtual human (similarly, part of the work by Gebhard et al. (2003) seeks to address the authoring problem by the design of better tools). An alternative may be to explore more automated approaches, a topic that we will take up in the next section.

5 TLTS and Thespian

The Tactical Language Training System (TLTS) (Johnson et al., 2004) is designed to teach the basics of a foreign language and cultural awareness. Figure 3 depicts a TLTS story. The language and cultural skills the learner must acquire in TLTS are mission or task related. In particular, we will discuss a task where the learner takes on the role of a soldier who is assigned to conduct a civil affairs mission in a foreign (e.g. Pashto, Iraqi) town. TLTS uses a 3D virtual world built on top of the Unreal Engine. The human user navigates in the virtual world and must interact with the virtual characters using spoken language and gestures to perform their mission.

Figure 3 TLTS

The virtual characters are autonomous characters that are part of the Thespian (Si, Marsella and Pynadath, 2005a,b) multi-agent based drama system. Each character is a Thespian agent. An automated speech recogniser identifies the utterance and converts them into a dialogue act representation that the agent-characters in Thespian take as input. Output from the Thespian agents consists of similar dialogue acts that instruct virtual characters to speak and behave.

The stories in TLTS consist of multiple scenes. The Pashto language story begins as the user arrives outside of a Pashto village. Some children are playing nearby and come over to talk to the user as the vehicle arrives. The user's aim in the scene is to establish initial rapport with people in the village and talking to their children in a friendly manner will further that goal.

TLTS is by far the most pedagogically ambitious of the projects discussed here in terms of wanting to develop an application that was going to be put to use by large numbers of learners. This had numerous implications. The story world discussed here is only part of the overall system; there are also various practice and game-like modules that allow the learner to rehearse pronunciation, vocabulary, etc. More significant for our discussion here, TLTS had to provide sufficient content for the learner to gain experience in the rudiments of a language, a challenging learning task. In the Iraqi Arabic version of TLTS the story has seven scenes, in the Pashto version there are four scenes. The process of working through these scenes takes considerable time. The design of the story and the characters should ideally support and maintain user interaction and engagement over the span and content of these scenes. If a learner became disengaged, the application would fail to give the learner the full training experience.

5.1 Pedagogy and narrative entanglement

The mission scenarios that form the basis of the scenes in TLTS were formulated by subject matter experts based on what would constitute rudimentary expertise in the foreign language and culture that was being taught as well as the language skills required for the mission task. This was different from CBI, where there was pre-existing, well thought out pedagogical material to draw on in crafting the interactive experience.

What was being taught made it critical that the pedagogical and domain experts, people who knew the language and culture and how to teach it, drive the agent authoring and not the technologists. Although this goal was also part of CBI and SASO design efforts, it was more critical in TLTS. Agent designers that had knowledge of Pashto language and culture were unavailable. It was impractical to think that agent designers could simply pick up language expertise in short order or that Pashto experts would rapidly pick up agent-design expertise. This was especially so because the TLTS projects were on a tight design schedule.

The Thespian agents that realise the characters in TLTS are defined within the PsychSim multi-agent architecture (Marsella, Pynadath and Read, 2004). PsychSim agents have beliefs that include recursive beliefs about other agents, goals, preferences over those goals and actions for achieving the goals. They also have both reactive and deliberative policies that drive how they select actions to achieve their goals. This is all cast within a decision-theoretic framework whereby an agent seeks to maximise its utility, the degree to which its goals are satisfied, weighted by its goal preferences. Finally, PsychSim has a built-in automated fitting algorithm for determining goal preferences based on instances of desired behaviour (Pynadath and Marsella, 2004).

Thespian restricts PsychSim agents in several ways. Thespian encodes the character's personalities, pedagogical goals, and social/cultural norm behaviours all as goals. Policies by which actions are selected are restricted to forward projection (look-ahead search). Specifically, Thespian agents simply choose actions based on reasoning about the impact that alternative actions will have on their utility, as weighted by their goal preferences. Changing those goal preferences leads to different behaviours. Consequently, authoring Thespian agents for interactive dramas can be viewed as the problem of adjusting goal preferences so that the agents have desired behaviour in a scene. Thespian relies on PsychSim's fitting algorithm as the basis of a new approach to authoring, as discussed in greater detail below.

5.2 Pedagogically mediated involvement

TLTS mission environment has an art design that appears rooted largely in realism. The learner interacts with the system using spoken dialog. In these ways, TLTS and SASO appear quite similar on the surface. However, there are a range of important differences. The TLTS environment is an augmented reality designed to provide affordances to the learner that facilitate their interaction with the system. For example, an arrow appears above the character that the learner is talking to in order to avoid confusion, trust meters above characters indicate how successfully the learner is in establishing trust, etc. These affordances are designed to simplify interaction with the system, make achievement of learning objectives clearer as well as to avoid frustration. As in SASO, the learner plays the main protagonist. Unlike SASO, the learner can see their character. She even has control over the camera position unlike the fixed camera of SASO or the agent-controlled camera of CBI.

5.3 Design process

With TLTS, we made a fundamental change in how we approached the relation between the authoring of the story and the authoring of the agents. We decided to bridge the gap between script and the implementation of the agent characters. Based on our experiences in CBI and SASO, we assumed authors would be more comfortable and faster at writing traditional linear stories. So rather than try to educate writers about agent design (as was attempted in SASO), we sought to make agent design more like script writing and less like programming.

Specifically, Thespian assumes that the starting point for the design process is a standard script or story outline, with possible variations, produced by a writer. The scripts provide 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 that adjusts agents' goals so that they are motivated to perform their roles according to the scripts. This ensures that the agents' autonomous behaviour can follow the script when the learner's behaviour is consistent with it, but is still true to their character's motivations even when the drama deviates from the script.

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. The dialog and actions in these scripts must be translated into a simple dialog and act formalism that Thespian can interpret. Once Thespian has fit the character's goals to this input, the character will always generate autonomous behaviour 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 behaviour over all situations. It is as if we were 'teaching' the agent the motivations of its character, as opposed to having them simply memorise the scripts. To make this process work, however, the story world must first be set up, including what objects are in the world, what actions the characters can take and how those actions impact the world. Currently, simplifying this set up of the story world is a focus of the research.

5.4 Thespian/TLTS discussion

Ideally, the Thespian authoring process reduces authoring effort. It 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 programming. The approach to authoring we have taken in Thespian is to bridge the art-technology gap by using automation to transform the design process so that it is closer to writing. To date, the approach appears promising. It speeds up development but technologists still must still be involved at critical junctures. Editing tools that would help authors translate their script into dialog acts and set up the story world are needed before we can fully assess the approach. Based on our initial work with authors, our current view is that creation of author-friendly versions of these tools will be possible.

6 Discussion

CBI, SASO and TLTS incorporate considerable technology. However, the process of creating these IPDs was not simply a technological endeavour. It involved a mixture of artistic, technological and most importantly pedagogical choices. Although the processes grappled with the issue of how the technological and artistic choices served the specific pedagogy, there were several design concerns shared across these systems.

All three systems resolve potential conflicts between pedagogical goals, drama and interactivity by following an approach of strong narrative entanglement. The entanglement is achieved with agent designs that directly embody the pedagogy or are fitted to it. Because the pedagogy is social in nature, especially for SASO and CBI, this lead to agent designs that are computational models of the psycho-social theories that lay behind the pedagogy. Ideally, this ensures that the learner's experience in the environment is consistent with the pedagogy.

All three systems mediate the learner's involvement based on pedagogical goals and specifically seek to regulate affective response. However, the desired affective response differs. For example, CBI's artwork and low-immersive approach to interaction design seeks to create psychological distance that reduces learner's personal stress. However, the evocative story seeks to promote empathy for Carmen. SASO's harsh realism and direct immersion is designed to create stressful emotions that the learner might face in the real world when dealing with similar challenges. TLTS on the other hand uses augmented reality to reduce learner frustration.

There were also similarities in the process of crafting the dramas. It always involved working from the professionally scripts that were essentially deconstructed to inform agent models. This deconstruction involved interpreting character motivation in terms of the pedagogy and incorporating that interpretation into the behavioural models of the agents.

The artefacts that result from these similar design processes are nevertheless fundamentally different. The CBI characters are the dramatic embodiment of two poles – one pole, Carmen, constitutes unhealthy emotional/coping reactions and on the other pole, Gina, constitutes the more deliberative, healthy ways to resolve problems and achieve more positive emotional states. Thus Carmen was largely driven by a psychological theory of stress while Gina was driven by a clinical model of healthy problem solving, her model of Bright IDEAS. SASO is a more open-ended negotiation with a flexible virtual human that is far more autonomous, with a deeper understanding of the content of the dialog it is engaged in and how it impacts its goals. Nevertheless, similar to CBI, the model that drives the doctor's negotiations is a psycho-social model that relates negotiation stances to emotional responses. The pedagogy in TLTS concerns language skills and appropriate, polite behaviour in a foreign culture. Although the agents incorporate models of how improper behaviour can influence such factors as trust, they do not possess deep models of emotions, their task or problem solving strategies as CBI and SASO characters do. What was required for TLTS was a more rapid way to move considerable story content, multiple scenes and characters, into the interactive game experiences. That led to exploring more automated (and less structured) ways to design a character's motivations so that it fit the pedagogy.

In trying to reconcile the demands of the pedagogy, art and technology, several lessons have been learned. We faced challenges in trying to get authors of traditional literature to appreciate the capabilities of this new media of agent-based interactive drama. The authors did best crafting linear stories or variations of linear stories with, for example, alternative outcomes. Even game designers, although familiar with the simpler AI-like technologies used in commercial games, may not easily understand the flexibility of an agent design well enough to inform how they can incorporate pedagogy, interactivity and variability into their stories. Also, traditional writers differ considerably in how readily their stories can inform agent design. Some authors write well-motivated characters that can quite readily be adapted to the autonomous agent technology we use. Other authors wrote stories that focused more on the plot and had character motivations that were less developed. Such stories require considerable interpretation to derive the character motivations. The design process flowed best when the pedagogy was well fleshed out beforehand. As we saw with CBI, the well-formed pre-existing pedagogy informed every aspect of agent design. However, this is not always feasible. In both SASO and TLTS, defining the pedagogy was part of the overall design process.

Finally, what is perhaps most apparent is that the space of alternative design approaches is large. Even approaches that seem similar on the surface, such as the three discussed in this article, on closer examination have numerous differences. It is the author's opinion that a large design space will always be faced when crafting interactive pedagogical dramas. Certainly there are mistakes to be made and lessons learned in avoiding parts of this design space. Nevertheless, the large design space is inherent in this blend of art, technology and pedagogy.

Acknowledgements

All the systems discussed here were collaborative projects. I would like to thank all my colleagues on these projects. I would especially thank W. Lewis Johnson and Catherine LaBore on CBI. On SASO, I would like to thank Bill Swartout, David Traum and Jonathan Gratch. On TLTS and Thespian, I would like to thank Mei Si, David Pynadath, Hannes Vilhjálmsson and W. Lewis Johnson. CBI was supported in part by the National Cancer Institute under grant # R25CA65520. TLTS is part of the DARWARS initiative of the Defense Advanced Research Projects Agency. SASO has been sponsored by the US Army Research Development and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government and no official endorsement should be inferred.

References

Bransford, J. and Vye, N.J. (1989) 'A perspective on cognitive research and its implications for instruction', in L. Resnick and L.E. Klopfer (Eds), *Toward the Thinking Curriculum: Current Cognitive Research*. Alexandria, VA: ASCD.

Bruner, J. (1990) *Acts of Meaning*. Cambridge, MA: Harvard University.

Cassell, J., Sullivan, J., Prevost, S. and Churchill, E. (Eds) (2000) *Embodied Conversational Agents*. Cambridge, MA: MIT Press.

Cavazza, M., Charles, F. and Mead, J.S. (2001) 'Agents' interaction in virtual storytelling', *Lecture Notes in Computer Science*, Vol. 2190, pp.156–170.

Cognition and Technology Group at Vanderbilt (1992) 'The Jasper series as an example of anchored instruction: theory, program description and assessment data', *Educational Psychologist*, Vol. 27, pp.291–315.

Galyean, T. (1995) *Narrative Guidance of Interactivity*. PhD thesis, Media Arts and Sciences, MIT, Cambridge, MA.

Gebhard, P., Kipp, M., Klesen, M. and Rist, T. (2003) 'Authoring scenes for adaptive, interactive performances', Paper presented in the Proceedings of the *International Joint Conference on Autonomous Agents and Multi-Agent System* (AAMAS'03) (pp.725–732). Melburne, Australia.

Gratch, J. and Marsella, S.C. (2004) 'A domain-independent framework for modeling emotion', *Journal of Cognitive Systems Research*, Vol. 5, pp.269–306.

Johnson, W.L., Marsella, S. and Vilhjalmsson, H. (2004) The DARWARS Tactical Language Training System, The Interservice/Industry Training, Simulation and Education Conference, 2004, Orlando, FL.

Johnson, W.L., Rickel, J. and Lester, J.C. (2000) 'Animated pedagogical agents: face-to-face interaction in interactive learning environments', *Int. J. AI in Education*, Vol. 11, pp.47–78.

Klein, G. (1998) *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press.

Lazarus, R. (1991) *Emotion and Adaptation*. New York, NY: Oxford University Press.

Lepper, M.R. and Henderlong, J. (2000) 'Turning play into work and work into play: 25 years of research in intrinsic versus extrinsic motivation', in C. Sansone and J.M. Harackiewicz (Eds), *Intrinsic and Extrinsic Motivation: the Search for Optimal Motivation and Performance*, (pp.257–307). San Diego: Academic Press.

Machado, I., Brna, P. and Paiva, A. (2006) 'Tell me a story', *Virtual Reality*, Vol. 9, pp.34–48.

Marsella, S., Gratch, J. and Rickel, J. (2004) 'Expressive behaviors for virtual worlds', in H. Prendinger and M. Ishizuka (Eds) *Life-like Characters Tools, Affective Functions and Applications*. Berlin: Springer-Verlag.

Marsella, S., Johnson, W.L. and LaBore, C. (2000) 'Interactive pedagogical drama', Paper presented in the Proceedings of the *Fourth International Conference on Autonomous Agents, Montreal, Canada*.

Marsella, S., Johnson, W.L. and LaBore, C. (2003) 'Interactive pedagogical drama for health interventions', Paper presented in the Proceedings of the *11th International Conference on Artificial Intelligence in Education (AIED), Sydney, Australia*.

Marsella, S., Pynadath, D.V. and Read, S.J. (2004) 'PsychSim: agent-based modeling of social interactions and influence', Paper presented in the Proceedings of the *International Conference on Cognitive Modeling*, pp.243–248.

Mateas, M. and Stern, A. (2003) 'Integrating plot, character and natural language processing in the interactive drama Facade', Paper presented in the Proceedings of the *International Conference on Technology for Interactive Digital Storytelling and Entertainment*.

McKee, R. (1997) *Story*. New York, NY: Harper Collins.

Onega, S. and Landa, J.A.G. (1996) *Narratology*. London: Longman.

Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L.E. and Zoll, C. (2005) 'Learning by feeling: evoking empathy with synthetic characters', *Applied Artificial Intelligence*, Vol. 19, pp.235–266.

Putnam, L.L. (1990) 'Reframing integrative and distributive bargaining: a process approach', in B.H. Sheppard, M.H. Bazerman and R.J. Lewicki (Eds), *Research on negotiation in organizations* (pp.3–30). Greenwich, CT: JAI Press.

Putnam, L.L. and Jones, T.S. (1982) 'Reciprocity in negotiations: An analysis of bargaining interaction', *Communications Monographs*, Vol. 49, pp.171–191.

Pynadath, D.V. and Marsella, S. (2004) 'Fitting and compilation of multi agent models through piecewise linear functions', Paper presented in the Proceedings of the *International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp.1197–1204.

Rasmussen, J.L. and Oakley, R.B. (1992) *Conflict Resolution in the Middle East: Simulating a Diplomatic Negotiation between Israel and Syria*. Washington, DC: United States Institute of Peace Press.

Read, S.J., Miller, L.C., Appleby, P.R., Nwosu, M.E., Reynaldo, S., Lauren, A. and Putcha, A. (2006) 'Socially optimized learning in a virtual environment: reducing risky sexual behavior among men who have sex with men', *Human Communication Research*, Vol.32, pp.1–34.

Rickel, J. and Johnson, W.L. (2000) 'Task-oriented collaboration with embodied agents in virtual worlds', in J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds), *Embodied Conversational Agents*. Boston: MIT Press.

Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D. and Swartout, W. (2002) 'Toward a new generation of virtual humans for interactive experiences', *IEEE Intelligent Systems*, Vol. 17, pp.32–38.

Riedl, M.O., Saretto, C.J. and Young, R.M. (2003) 'Managing interaction between users and agents in a multi-agent storytelling environment', Paper presented in the Proceedings of the *International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp.741–748.

Schank, R.C. (1996) 'Goal-based scenarios: case-based reasoning meets learning by doing', in D. Leake (Ed.), *Case-Based Reasoning: Experiences, Lessons & Future Directions* (pp. 295–347). Cambridge, MA: MIT Press.

Shaftel, F.R. and Shaftel, G. (1982) *Role Playing in the Curriculum*. Englewood Cliffs, NJ: Prentice-Hall.

Si, M., Marsella, S. and Pynadath, D.V. (2005a) 'Thespian: using multi-agent fitting to craft interactive drama', Paper presented in the Proceedings of the *International Conference on Autonomous Agents and Multi Agent Systems*, pp.21–28.

Si, M., Marsella, S. and Pynadath, D.V. (2005b) 'Thespian: an architecture for interactive pedagogical drama', International Conference on Artificial Intelligence in Education, Amsterdam.

Sillars, A.L., Coletti, S.F., Parry, D. and Rogers, M.A. (1982) 'Coding verbal conflict tactics: nonverbal and perceptual correlates of the avoidance-distributive-integrative distinction', *Human Communications Research*, Vol. 9, pp.83–95.

Swartout, W., Gratch, J., Hill, R., Hovy, E., Lindheim, R., Marsella, S. Rickel, J. and Traum, D. (2005) 'Simulation meets Hollywood: integrating graphics, sound, story and character for immersive simulation', in O. Stock and M. Zancanaro (Eds), *Multimodal Intelligent Information Presentation* (pp.279–303). Berlin: Springer.

Szilas, N. (2003) IDtension, a narrative engine for interactive drama. The 1st International Conference on Technologies for Interactive Digital Stroytelling and Entertainment, Darmstadt, Germany.

Traum, D. and Rickel, J. (2002) 'Embodied agents for multi-party dialogue in immersive virtual worlds', Paper presented in the Proceedings of the *1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pp.766–773.

Traum, D., Swartout, W., Marsella, S.C. and Gratch, J. (2005) 'Fight, flight or negotiate: believable strategies for conversing under crisis', Paper presented in the Proceedings of the *5th International Working Conference On Intelligent Virtual Agents September 12–14, Kos, Greece*.

Varni, J.W., Sahler, O.J., Katz, E.R., Mulhern, R.K., Copeland, D.R., Noll, R.B., Phipps, S., Dolgin, M.J. and Roghmann, K. (1999) 'Maternal problem-solving therapy in pediatric cancer', *Journal of Psychosocial Oncology*, Vol. 16, pp.41–71.

Vygotsky, L. (1978) 'Mind in society: the development of higher psychological processes', M. Cole, V. John-Steiner, S. Scribner and E. Souberman (Eds.), and Trans. Cambridge, UK: Cambridge University Press.

Walton, R.E. and Mckersie, R.B. (1965) *A Behavioral Theory of Labor Negotiations: an Analysis of a Social Interaction System*. New York, NY: McGraw-Hill.

Weyhrauch, P. (1993) *Guiding Interactive Drama*. PhD thesis, Carnegie Mellon University. Technical Report CMU-CS-97-109.

Wilson, S.R. and Putnam, L.L. (1990) 'Interaction goals in negotiation,' *Communication Yearbook*, Vol. 13, pp.374–406.

Yardley-Matwiejczuk, K.M. (1997) *Role Play: Theory and Practice*. London: Sage Publications Ltd.

Young, R.M. (2001) 'An overview of the Mimesis architecture: integrating intelligent narrative control into an existing gaming environment', The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment, Stanford, CA.