Intelligent Narrative Technologies: Papers from the 2012 AIIDE Workshop
AAAI Technical Report WS-12-14

Gestural Interactions for Interactive Narrative Co-Creation

Andreya Piplica, Chris DeLeon, and Brian Magerko
Georgia Institute of Technology
{apiplica3, cdeleon3, magerko}@gatech.edu

Abstract

This paper describes a gestural approach to interacting with interactive narrative characters that supports co-creativity. It describes our approach using a Microsoft Kinect to created a short scene with an intelligent avatar and an AI-controlled actor. It describes our preliminary user studies and a recommendation for future evaluation.

Introduction

One under-researched approach to interactive narrative seeks to allow humans and intelligent agents to *co-create* narratives in real-time as equal contributors. In an ideal narrative co-creation experience, a human interactor and an intelligent agent act as characters in a novel (i.e. not pre-scripted) narrative. The experience is valuable both as a means of creating a narrative and engaging in a performance. In such a system, the story elements for an experience are not pre-defined by the AI's knowledge (e.g. a drama manager who is imbued with the set of possible story events for an experience), but rather the human and AI have a) similar background knowledge and b) similar procedural knowledge for how to collaboratively create a story together. While this stance is within the boundaries of *emergent narrative* (i.e. approaches to interactive narrative that do not have prescribed story events), the typical approach in emergent narrative systems relies on strongly autonomous systems (i.e. rational agents pursuing their goals to elicit story events), which may not engage directly with a user and their narrative goals. Co-creation describes a more collaborative process between a human interactor and an AI agent or agents that relies on *equal contributions* to the developing story. This article focuses on one of the major problems with co-creative systems and interactive narratives in general: how the human user interacts with the story world in an immersive fashion.

While intelligent agents in interactive narratives can interact with humans through expressive embodied representations, human interactors have no way to reciprocate with embodied communication of their own. Embodied agents can move, portray facial expressions, and convey dialogue. Humans, on the other hand, are typically limited to interacting through mouse and keyboard, which do not afford much physical expression. This disparity in interaction capabilities has been informally discussed in the community as the *human puppet problem*.

The two main paradigms for human interaction in interactive narratives are language-based and menu-based interaction. With language-based interaction, users type or speak natural language utterances to interact with virtual characters (Mateas and Stern 2003; Riedl and Stern 2006; Aylett et al. 2005). With menu-based interaction, users select from dialogue or actions in a visual list (McCoy et al. 2011). Some interactive narratives respond to the user's actions in a game world, which are still mediated by keyboard-and-mouse input modalities (Magerko and Laird 2003; Young 2001). None of these approaches involves full-bodied gestural interaction directly from the user. Embodied agents can add animated gestures to a user's input (Zhang et al. 2007; El-Nasr et al. 2011), but this does not allow the user to communicate with their own gestures. Fully immersive environments (Hartholt, Gratch, and Weiss 2009) allow users to interact with virtual characters and objects as they would real people and objects, but often focus on established stories rather than co-creating new ones.

Improvisational theatre (improv) provides a real-world analogue to embodied co-creative interactive narrative experiences (J. Tanenbaum and Tanenbaum 2008; Magerko and Riedl 2008). Improvisational actors collaboratively create novel stories in real-time as part of a performance. They communicate with each other through both dialogue and full-body movements. Movements can portray characters (including their status, mood, and intent) (Laban and Ullmann 1971) as well as contribute actions to a scene which establish the activity or location (Johnstone 1999). All communications occur within the (diegetic) context of an improvised scene; therefore, improvisers cannot use explicit communication to resolve conflicting ideas about the scene they are creating (Magerko et al. 2009). In an ideal real-time co-creative system, human interactors and AI agents communicate with each other through their performance rather than through behind-the-scenes communication. Thus, the physical and diegetic qualities of improv make it an ideal domain to study in order to inform technological approaches to embodied co-creation.

This paper presents a system for combining improv acting with full-body motions to support human-AI co-creation of interactive narratives. Our system translates a human interactor's movements into actions that an AI improviser can respond to in the context of an emergent narrative. We present a framework for human interaction with an AI improviser for beginning an improvised narrative with interaction mediated through an intelligent user avatar. This approach uses human gestural input to contribute part, but not all, of an intelligent avatar's behavior. We also present a case study of trained improvisers interacting with the system, which informs challenges and design goals for embodied co-creation. While joint human-AI agents have been employed in interactive narrative (Hayes-Roth, Brownston, and van Gent 1995; Brisson, Dias, and Paiva 2007), this is the first system to do so with full-body gestures.

Gestural Narrative Co-creation

Interactive narratives have used embodied gestural interaction in stories with established characters (Cavazza et al. 2004) or an established dramatic arc (Dow et al. 2007), though not for narrative co-creation. We enable gestural interaction in our architecture, with co-creativity as a goal, by combining elements from two real world improv games, *Three Line Scene* and *Moving Bodies* (Magerko, Dohogne, and DeLeon, 2011). In *Three Line Scene*, actors establish the *platform* of a scene (i.e. the scene's characters, their relationship, their location, and the joint activity they are participating in (Sawyer 2003)) in three lines of dialogue. Improvisers establish the platform at the beginning of a scene to create context before exploring a narrative arc. We focus only on the character and joint activity aspects of a platform as a simplification of the knowledge space involved, since location and relationship are often established implicitly or omitted (Sawyer 2003). In *Moving Bodies*, improvisers provide dialogue for a scene as usual, but let other people (typically audience members) provide their movements. The mover poses the improviser as if controlling a puppet while the improviser creates dialogue based on their interpretations of these poses.

In our system, the human interactor and an AI improviser improvise a pantomimed three-line scene. An intelligent avatar uses motion data from a Microsoft Kinect sensor to represent the human's motions in the same virtual space as the AI improviser. It gives the human interactor an embodied presence in the scene and shows how the Kinect senses their motion. This feedback can help the user understand the avatar's interpretations and adjust their movements to accommodate the sensor's limitations. The avatar reasons about the human's intentions for the scene and creates its own mental model. Future implementations of the system will follow the *Moving Bodies* interaction metaphor with both the AI improviser and the intelligent avatar producing dialogue based on their models of the scene. The current implementation omits dialogue in favor of studying gestural interaction with improvisational agents alone.

The user stands before a large screen where the AI improviser and the intelligent avatar are displayed on a stage (Figure 1). The user faces the screen while standing about four to ten feet away from a Microsoft Kinect below the screen (i.e. within the Kinect's sensor range). Both the intelligent avatar and the AI improviser are shown as two-dimensional animated characters. These simplified representations map motion data from the Kinect directly onto the characters' animations. The system does not currently support animated facial expressions, though such animations may be supported in future iterations.

The user performs a motion to begin a scene, such as putting one fist on top of the other and moving their hands from side to side. The Kinect sends the motion data to the intelligent avatar and the AI improviser, who each interpret the motion as an action. The avatar displays the user's motion while it reasons about what character and joint activity the user may be portraying. Here, the avatar may identify the user's motion as, for example, the action *sweeping* and reason that they are portraying the character *shopkeeper*. The AI improviser reasons about the user's

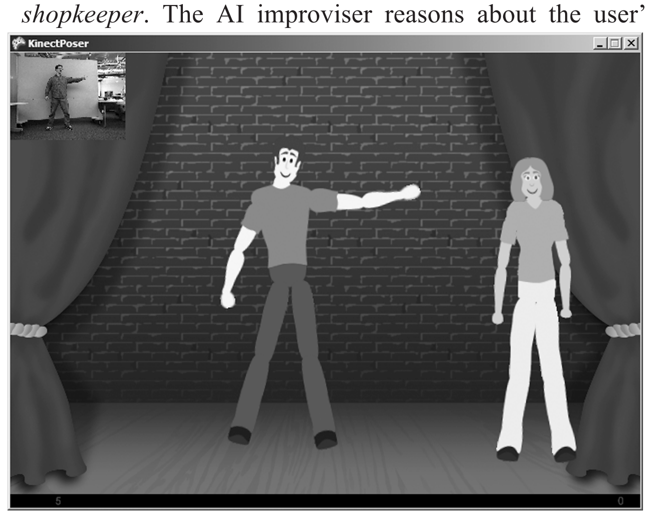

Figure 1. Human (upper left) controlling the intelligent avatar (middle) while performing a three-line scene with an AI improviser (right).

character and joint activity as well as its own character. Here it may reason that the user is *sweeping* as part of the joint activity *cleaning*. The AI improviser then chooses an action to present (e.g. *wiping tables* as part of the joint activity *cleaning*). It presents this motion to the user through an animation. This animation would show the improviser moving one hand at waist height repeatedly in a circle. Then it is the user's turn to present another motion to the scene.

AI-Based Improviser and Avatar

The AI improviser and the intelligent avatar draw on the same reasoning processes to understand how the human interactor contributes to the scene. They utilize background knowledge about a specific domain to make inferences about the platform. They incorporate the human's motions into their reasoning based on joint position data from the Microsoft Kinect sensor. Additionally, the AI improviser reasons about how to contribute presentations to the scene with motions. The intelligence given to the intelligent avatar is intended to provide additional language capabilities in future iterations of the architecture. If the avatar can make reasonable interpretations of user gestural input based on a) background knowledge and b) context given from the scene, then it can add reasonable dialogue utterances much like an improviser would that is performing the *Moving Bodies* improv game.

Background Knowledge

Both the AI improviser and the intelligent avatar reason about how an improviser (human or AI) contributes to a scene based on the fuzzy inference agent framework described in (Magerko, Dohogne, DeLeon 2011). This framework describes how agents can reason about characters and joint activities in a scene based on motion and dialogue inputs to infer fuzzy concepts and produce perfomative outputs. While both the avatar and AI improviser rely on procedural definitions of how to negotiate a platform, they also need background knowledge about a story world for those procedures to operate on. The contents of this background knowledge base can refer to any narrative domain; we have confined our narrative domain, called *TinyWest*, to a set of actions, characters, and joint activities associated with stories in the Old West. The elements of the Old West genre have been codified into generally recognizable stereotypes so that we can assume that most people have a similar understanding of these elements (e.g. most people have similar ideas about what a cowboy is and does).

We collected crowdsourced data through Amazon's Mechanical Turk to find degree of association values (i.e. bi-directional fuzzy relations) (O'Neill, et al. 2011) relating 59 actions, 11 characters, and 12 joint activities. We chose the actions, characters, and joint activities based on analysis of sixteen Western genre films and television shows. Mechanical Turk allows us to survey a large number of people and gather multiple data points for each pair of associated elements. We can represent real-world variability in human opinions by assigning the AI improviser and the avatar different values for the same pair of associated elements. Forcing the agent and avatar to have slightly different background knowledge prevents them from agreeing on aspects of the scene automatically, making their interactions more like two improvisers negotiating their mental models.

Interpreting Motions as Actions

When the human interactor presents a motion to the scene, both the intelligent avatar and the AI improviser need to interpret the Kinect data for the motion as a semantic action. Motions do not contain semantic information on their own; the same motion presented in different contexts can portray several actions (Kendon 2004). For example, if an actor puts their fists together, brings them up to their shoulders, and swings them in front of them, they could be pantomiming swinging a baseball bat or chopping a tree with an axe, depending on the scene.

The motion data from the Microsoft Kinect consists of 3D coordinates for joint and limb positions, which we project to 2D and use to animate the on-screen characters. The coordinates are evaluated as "signals," which are mathematical representations of specific joint angles, relative joint positions, or changes over time. Signals can be *simple* (joint angles and positions at one time) or *temporal* (joint angles and positions varying across time). For example, the *arms crossed* simple signal detects whether each hand is positioned close to the other arm's elbow. The *punching* temporal signal, in contrast, detects whether either arm rapidly extends and returns. Performing this motion too slowly does not trigger the *punching* signal. Temporal signals are evaluated over the span of a user's turn. A turn is considered complete when the human interactor has either been out of a neutral stance for a set period of time (currently two seconds) or returns to a neutral stance after the system identifies at least one candidate action. A neutral stance is defined as standing still with feet together and hands at one's sides.

Actions are defined as sets of positive and negative signals. If the Kinect data satisfies all positive signals for an action, it becomes a candidate. However, if the data triggers any negative signals for that action, it is removed as a candidate. When the user's turn ends, the intelligent agent and AI improviser select an action from the candidate interpretations based on their mental model of the scene (O'neill, et. al. 2011).

Deciding on a Motion to Present

After the AI improviser interprets the user's motion and reasons about how it contributes to the scene, the AI improviser must select an action to present and a motion to present it with. The AI improviser selects an action to present in a similar way to how it reasons about characters and joint activities (O'neill, et al., 2011). Once the AI improviser selects a motion, its on-screen character plays the corresponding animation. In our initial authoring, we assigned one motion animation to each action. We recorded motion data from the Kinect of one author performing his interpretation of each motion. The author posed in several successive keyframe poses for each motion while the Kinect captured his body position at three second intervals. Posing for keyframes rather than recording continuous movements gave the author greater control over which frames of motion the Kinect captured. We took this approach to animation authoring for two reasons: 1) by using an animation system of our own, rather than an outside tool, we were able to test and iterate on the signals and gesture recognition functionality designed for the user character, and 2) this approach simplified and accelerated animation work, which was necessary since many animations were required and the authors are not trained in character animation.

User Studies

Our system evaluation focused on case studies of how expert improvisers interact with prototypes of our AI improviser and intelligent avatar. We studied improvisers from a local improv troupe as an initial case study of how gesture can be used as a main interface paradigm for interactive narrative.

We observed human improvisers playing our pantomime version of *Three-Line Scene*, where the human and AI improvisers attempt to establish the platform entirely through pantomime. Our recruitment with a local theatre yielded four improviser participants (one female, three male) . Instead of ending the scene after three lines of dialogue, the scenes ended after three motion exchanges between the human and AI improvisers. This resulted in six motions per scene, three from the human and three from the AI improviser. The human contributed the first motion to each scene and waited for the AI to respond. The avatar still processed the human improviser's motions and reasoned about actions, characters, and joint activities.

After a practice scene to introduce them to the system, each human improviser performed two scenes with the AI improviser in the *Tiny West* domain. In the first scene, the human received no feedback regarding the AI's interpretation of their motion. In the second, the system displayed a one or two word description of the action the AI interpreted the improviser's motion as. These two conditions helped us study how displaying the AI improviser's interpretations of the user's movements affects the user's understanding of and engagement with the scene. The improvisers filled out a questionnaire about their impressions after each scene. Part way through the questionnaire, we revealed the AI improviser's interpretation of characters and joint activity in the scene and asked the human to evaluate these interpretations. After the second questionnaire, we conducted a brief, structured interview with the participant.

Interpretation Feedback. The improvisers responded to the on-screen interpretations with mixed feedback. Two improvisers found that seeing the interpretation helped them direct the scene. Improviser 3 considered changing what he was doing to align with the AI's interpretation. Improviser 2 rated the interpretations as helpful (4 on a Likert scale from 1 to 5, where 1 is "not at all helpful" and 5 is "extremely helpful"). This improviser thought that the AI's models of the scenes were at least moderately reasonable in both scenes (at least 3 on a Likert scale from 1 to 5, where 1 is "completely unreasonable (seems random)" and 5 is "completely reasonable (interpretations make sense)"), the only improviser to do so.

In contrast, the displayed interpretations sometimes confused the human improvisers, especially when the interpretations diverged from their own. When Improviser 3 encountered a divergent interpretation, he said it "took me out of the process of building a scene." He described it as "jarring that they're putting a label, and that's not at all what I thought." Improviser 1 reported feeling unsure how to react upon seeing that the AI improviser's interpretation of her motion differed from her own. If she encountered such a divergence in a scene with another human improviser, she would have accepted their interpretation and built off of that, a practice called "yes-and"-ing (Johnstone 1999). Instead, she chose to follow her original interpretation. This fits well with our understanding of how ambiguity is used on stage (Magerko, Dohogne, and DeLeon, 2011); improvisers rely on ambiguity in a performance to allow for multiple possible interpretations of the scene elements as they work to create a shared mental model.

Improvisers rated the reasonableness of the AI improviser's interpretations higher when they received no feedback about how it interpreted their motions. That is, when the human improvisers judged the AI's interpretations of their motions only by observing its motions, the humans considered the AI's interpretations more reasonable than when they saw explicit descriptions of its interpretations. Three out of four human improvisers rated the AI improviser's mental model interpretations as random (1 on a Likert scale from 1 to 5, where 1 is "completely unreasonable (seems random)" and 5 is "completely reasonable (interpretations make sense)").

AI Improviser Movements. Three out of the four improvisers reported difficulties observing and understanding the AI improviser's motions. Improviser 1 reported, "I don't know what happened on their [the agent's] part," claiming that the agent's motions looked like "just a bunch of movements." Improviser 4 described the agent's movements as "random," "undefined," and "difficult to interpret." Improvisers 1 and 2 both described the AI improviser's motions as "fast." They may have been referring to the animations being short in duration or the AI improviser presenting while they were still performing. Improviser 2 said, "Normally, I'd let things build a little bit to kind of feel out where I had a scene," indicating that he was still in the process of performing motions for his turn when the agent presented. (Like the improvisers in the previous study, these improvisers typically performed multiple discrete motions in one turn.) Improviser 4 elaborated on the difficulties of turn-taking, saying, "We either need to be able to engage simultaneously or engage very, very distinctly. This felt like it was neither of those." Improviser 1 suggested that letting the AI start the scene would ease the turn-taking at the beginning since the human is better equipped to interpret motions and add to a narrative than the artificial system. This result may point to the need to work with improvisers in the future who are trained in improvisational pantomime and the need to jointly develop a gestural discourse vocabulary.

Differences from Improvising with Humans. The improvisers noted a marked difference between interacting with this system and their traditional improv experience. Improviser 1, as noted above, seemed unsure whether to interact with the AI as she would with a human improviser. She was unsure how well the AI improviser would be able to interpret her motions, saying, "I don't know that it's intelligent enough to really know what I'm doing." Improviser 3 expressed a similar uncertainty and suggested letting interactors "test ahead of time different gestures to see what the system recognizes." The human improvisers could not treat the agent as an equal in their improvisations because they could not assume a particular level of shared knowledge. Additionally, the humans still felt like they lacked non-verbal communication modalities even though they could communicate with gestures. Improviser 4 specifically mentioned a lack of eye contact, which he could use with another human improviser to establish that they were both "on the same page."

Preference for Face-to-Face Interaction. We initially thought that displaying the improvisers' actions on-screen with the avatar would help them understand how the AI improviser interpreted their movements. Three out of the four improvisers said they would prefer a face-to-face interaction with the agent rather than the "stage-view" that showed both the AI improviser and the improviser's avatar. Improviser 1 missed presentations from the AI improviser because she attended too much to the avatar. A face-to-face display would immerse improvisers in the scene more since the display would more closely mirror their on-stage experience. Improviser 3, however, felt that the stage-view "puts you in the scene." He found that seeing himself moving was helpful in a way that seeing "just what your character sees" would not be.

Discussion

Our studies of expert improvisers interacting with improvisational agents through full-body gestures indicate several goals and challenges for future interaction designs, encompassing both the AI improviser's background knowledge and the way that the scene is displayed to the interactor. (The small sample size of our case study prevents us from making broader generalizations.) Balancing an improviser's traditional experience with the need to present clear information about the AI's contribution to the scene presents the main challenge of designing embodied co-creative interaction with intelligent agents for improvisers. While the visual display should mimic traditional improv as much as possible to create an immersive experience, other aspects of interaction need to accommodate the AI improviser's shortcomings to support the human's agency and clarity in the scene. Designing experiences for people without improvisation training should focus more on clearer presentations than faithfulness to traditional improv. A visual display that mirrors a traditional improvisation experience as much as possible will create the best sense of immersion for a human improviser. The screen should not show interpretations of the improviser's movements, since seeing explicit, non-diegetic interpretations of divergences reduced the human's sense of immersion. Not having this feedback will help human improvisers treat the AI improviser more like a fellow human improviser. If the human improviser is more comfortable with the AI improviser, their interactions will feel more natural and immersive. While we initially thought giving the user an embodied presence in the same virtual space as the AI improviser would make interaction more understandable, this full stage view differs too much from the improviser's typical experience. Interaction with the AI improviser should be presented as face-to-face, as improvisers would interact this way on-stage. A transparent outline of the human improviser's motions in the foreground may be necessary to represent the human in the virtual space in future work when the intelligent avatar produces dialogue on the human's behalf as in *Moving Bodies*.

The human improvisers' perception that their interactions with the AI improviser were random arose from unclear presentations from the AI. This perceived randomness in turn inhibited the human improvisers' sense of agency because they felt that their presentations did not cause a sensible response from the AI improviser. While ambiguity and unexpected presentations are a natural part of improv, presentations should still make sense within the context of a scene. One factor contributing to this may have been the lack of more subtle non-verbal communication from the AI, for example in the form of facial expressions, eye movements, or small changes in posture, that could provide clues to both the AI's interpretation of the improviser's action and the intention in responding. Improving the quality of the AI improviser's motion animations so that they are as life-like as possible will alleviate the perception of randomness and improve both clarity and agency. Chained and repeated motions in a human improviser's presentations made it harder for the AI improviser to moderate turn-taking. To ease this difficulty and to increase the human improviser's sense of agency, the user should explicitly indicate when their turn ends. The human improviser will clearly know when to attend to the AI improviser. An invented hand signal (which the improviser will not use in their natural presentations) can communicate the end of a turn. Although human-moderated turn-taking departs from the continuous flow of presentations in traditional improv, this sacrifice enhances the clarity of presentations. These results indicate the need for studies in the future that remove the AI from the system and rely on a Wizard of Oz technique to clearly focus on the gestural interactions without distraction from how intelligent the AI may or may not be.

References

Aylett, R. S., S. Louchart, J. Dias, A. Paiva, and M. Vala. 2005. FearNot!-an experiment in emergent narrative. In *Intelligent Virtual Agents*, 305–316.

Brisson, A., J. Dias, and A. Paiva. 2007. From Chinese Shadows to Interactive Shadows: Building a storytelling application with autonomous shadows. In *Proceedings of the Agent-based Systems for Human Learning Workshop (ABSHL) at Autonomous Agents and Multiagent Systems (AAMAS)*, 15–21.

Cavazza, M., F. Charles, S.J. Mead, O. Martin, X. Marichal, and A. Nandi. 2004. "Multimodal acting in mixed reality interactive storytelling." *Multimedia, IEEE* 11 (3): 30–39.

Dow, S., M. Mehta, E. Harmon, B. MacIntyre, and M. Mateas. 2007. Presence and engagement in an interactive drama. In *CHI-CONFERENCE-*, 2:1475.

El-Nasr, M., K. Isbister, J. Ventrella, B. Aghabeigi, C. Hash, M. Erfani, J. Morie, and L. Bishko. 2011. "Body buddies: social signaling through puppeteering." *Virtual and Mixed Reality-Systems and Applications*: 279–288.

Hartholt, A., J. Gratch, and L. Weiss. 2009. At the virtual frontier: Introducing Gunslinger, a multi-character, mixed-reality, story-driven experience. In *Intelligent Virtual Agents*, 500–501.

Hayes-Roth, B., L. Brownston, and R. van Gent. 1995. Multiagent collaboration in directed improvisation. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, 148–154.

Johnstone, K. 1999. *Impro for storytellers*. Routledge/Theatre Arts Books.

Laban, R., and L. Ullmann. 1971. "The mastery of movement."

Magerko, B., and J. E. Laird. 2003. Building an interactive drama architecture. In *First International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 226–237.

Magerko, B., and M. Riedl. 2008. "What Happens Next?: Toward an Empirical Investigation of Improvisational Theatre." *5th International Joint Workshop on Computational Creativity*.

Magerko, B., W. Manzoul, M. Riedl, A. Baumer, D. Fuller, K. Luther, and C. Pearce. 2009. An Empirical Study of Cognition and Theatrical Improvisation. In *Proceeding of the Seventh ACM Conference on Creativity and Cognition*, 117–126.

Magerko, B., Dohogne, P. and DeLeon, C. 2011. Employing Fuzzy Concepts for Digital Improvisational Theatre. In the Proceedings of the Seventh Annual AI and Digital Interactive Entertainment Conference, Palo Alto, CA.

Mateas, M., and A. Stern. 2003. Integrating plot, character and natural language processing in the interactive drama Façade. In *Proceedings of the 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE-03)*.

McCoy, J., M. Treanor, B. Samuel, N. Wardrip, and M. Mateas. 2011. Comme il faut: A system for authoring playable social models. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

O'Neill, B., Piplica, A., Fuller, D., and Magerko, B. 2011. A Knowledge-Based Framework for the Collaborative Improvisation of Scene Introductions. In the Proceedings of the 4th International Conference on Interactive Digital Storytelling, Vancouver, Canada.

Riedl, M., and A. Stern. 2006. "Believable agents and intelligent story adaptation for interactive storytelling." *Technologies for Interactive Digital Storytelling and Entertainment*: 1–12.

Sawyer, R. K. 2003. *Improvised dialogues: Emergence and creativity in conversation*. Ablex Publishing Corporation.

Tanenbaum, J., and K. Tanenbaum. 2008. "Improvisation and performance as models for interacting with stories." *Interactive Storytelling*: 250–263.

Young, R.M. 2001. An overview of the mimesis architecture: Integrating intelligent narrative control into an existing gaming environment. In *The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, 78–81.

Zhang, L., M. Gillies, J. Barnden, R. Hendley, M. Lee, and A. Wallington. 2007. "Affect detection and an automated improvisational ai actor in e-drama." *Artifical Intelligence for Human Computing*: 339–358.