Viewpoints AI

Mikhail Jacob, Gaëtan Coisne, Akshay Gupta, Ivan Sysoev, Gaurav Gav Verma, Brian Magerko

Georgia Institute of Technology

{mikhail.jacob, gcoisne3, akshaygupta, ivan.sysoyev, verma.gav, magerko}@gatech.edu

Abstract

This article describes a technical approach to movement-based interactions between a human interactor and an intelligent agent based on the theatrical Viewpoints movement framework. The *Viewpoints AI* system features procedural gesture interpretation using shallow semantics and deep aesthetics analysis from the Viewpoints framework. The installation creates a liminal virtual / real space for the human and AI to interact by the use of digital projection for the AI visualization and shadow play to represent the human. Observations from a recent public demonstration of the system and future directions of work are also discussed.

Introduction

Building interactive narrative experiences is a highly difficult challenge from both an artistic and technical perspective that is typically tackled with techniques like drama management, story generation, and intelligent virtual actors (Riedl and Bulitko 2013). These works all tend to deal with the infamous "authoring bottleneck," which refers to the difficulty in authoring content for interactive experiences. In other words, the more agency or meaningful choices given to the user, the more difficult it is to author content (and connections between the content) that cover the myriad narrative paths the user may experience and so procedural content generation affords higher user agency. However, instantial (non-procedural) content (e.g. story beats, typical virtual environments, etc.) can be less generic since it can be tailored to the desired user experience.

Our work presented here addresses an explicit design challenge: what kind of interactive narrative experience can be offered with a high degree of user agency and low authoring effort? In other words, what kind of experience can we create with only minimal, if any, instantial content? Tackling such a challenge suggests specific constraints and approaches: 1) non-verbal language interaction should not be used (far too cumbersome for a low-authoring approach); 2) the absence of verbal interaction infers that story construction should be based on movement and gesture; and 3) the use of movement and gesture need not to be tied directly to deep semantics, else the same problem exists with motion as does with verbal interactions in terms of comprehension, building specific storylines, etc.

Addressing the previous design challenge has pushed us to consider how to create interactive protonarratives (e.g. highly abstract, semantically ambiguous narratives) that are created by the collaborative movement of a human user and intelligent agent. We have worked with a local contemporary theatre group to use contemporary movement theory in storytelling as a means of inspiring a new procedural approach to creating interactive narrative / artworks, called *Viewpoints AI*. We have employed the Viewpoints movement theory (see Section Viewpoints) as a formalism that is detailed enough to provide protonarratives with surface semantic reasoning while not concerning itself with being too abstract or too specific, and thus bringing up issues of deep semantic understanding.

The *Viewpoints AI* system has been created as an exploratory AI / art research platform for understanding how

Figure 1: A human interactor and VAI co-creating a movement-based proto-narrative in a liminal virtual / real projected plane.

the aforementioned design challenge and constraints can be approached. It is an installation piece that involves a human user doing improvisational shadow play with an ephemeral intelligent agent that is projected onto the same 2D space as the user's shadow (see Figure 1). *Viewpoints AI* perceives the human's gesture, defined here as body movement between two consecutive periods of stillness; analyses it broadly through the Viewpoints perspectives of *space*, *time*, *emotion*, *shape* and *movement*; and finally responds to it generating an aesthetically appropriate response based on the human's gesture and Viewpoints technique. The resulting experience is a highly interactive, movement-based experience that does not rely on pre-authored instantial content, but instead on the procedural definition of Viewpoints theatrical movement theory and the user's own movements.

The *Viewpoints AI* system also provides the user with a liminal (i.e. in between the virtual and real) co-creational space, based on shadow play, that allows for a high degree of presence. This liminal space is created by front-projecting (from the audience's perspective) *VAI*'s virtual avatar onto a semi-opaque muslin screen. The user's shadow is rear-projected (from the audience's perspective) onto the same screen in order to allow a human user to interact with *VAI* through shadow play. The 2D plane of interaction serves as a creative equalizer for *VAI* and the human. The human's shadow serves as a particularly crisp digital avatar, retaining analog fidelity, while being able to interact with a digitally projected virtual interactor.

Related Work

The use of computation to generate creative and interactive productions is on the rise, with several pieces featured in popular forms of media such as games, as well as in art galleries (McCorduck 1991; Mateas 2003; Colton 2011). For example, *Spore* made heavy use of procedural content generation (PCG) to generate skin meshes, textures and animations on anthropomorphic creatures morphed or created by users (Hecker et al. 2008).

Researchers have explored the nature of creative composition and how it relates to the development of interactive computer-based tools to assist the composer, specifically for dance and movement. *Life Forms* (Calvert et al. 1991) is a computational tool that assists dance composers by utilizing spatio-temporal representations of the emerging composition. *EMVIZ* (Subyen et al. 2011) visualizes the Laban Basic Efforts (Laban et al. 1975) as an aid to appreciating the qualities of movement.

Related research on the usage of theatrical frameworks and techniques for creating entertainment agents exists in domains ranging from games to interactive narrative. The *Improv* system (Perlin 1996) used scripting and rules to imbue virtual actors with more realistic behaviors for virtual scenes. The *Digital Improv Project* (Magerko 2009; O'Neill 2011) modeled the cognitive processes occurring within improvisational actors such as shared mental model negotiation, platform creation and iconicity in order to create theatrical improvisational agents to play the improv games *Three Line Scene* and *Party Quirks*. The *Digital Improv Project* used gestural interaction in order to co-create narrative in the Wild West domain but was restricted to instantial gestures that triggered specific actions that the user could do to proceed with the narrative (Piplica 2012).

El-Nasr used Stanislavskian methods for embodying agents with motivation and goals in the adaptive interactive narrative *Mirage* (El-Nasr, 2007). Stanislavskian methods involve reasoning about character motivations and using different tactics to achieve these motivations and goals. Sharma et al. used the Laban Effort System to author flight paths for a quadrotor in order to induce affective responses in viewers (Sharma et al. 2013).

Sensing and projection have also been used to augment movement and dance performances such as in the piece *The Angled Angels Assembly* (Latulipe et al. 2011). In *The Angled Angels Assembly*, an overhead camera tracked dancers in order to visualize their movements using a procedural visualization projected on stage. Sensing and projection were also used in the theater play *It/I* (Pinhanez and Bobick 2002), where a projected digital character (It) responded to specific poses of a human actor (I) on a digitally augmented stage to act out a script together.

Since early Imperial China, shadow theater has been used in a variety of contexts including political commentary, rituals, religion, art and entertainment (Chen 2007). Shadow theater has seen resurgence in popular culture, for example on *America's Got Talent* (Cowell et al. 2006) entries such as *Team Attraction* and *Team iLuminate*. Traditional Greek shadow theater techniques have also been used in *EShadow* (Christoulakis et al. 2013), a tool for authoring stories in a virtual, collaborative environment.

Viewpoints

Viewpoints is a compositional technique for dancers to communicate through movement and gesture, building a vocabulary based on motions and allowing dancers to improvise in a defined environment. It was developed in the 1970s by Mary Overlie (2006) and was later adapted for stage acting by Anne Bogart and Tina Landau (2005). It is a tool for developing spontaneous interactions between actors and is also used to train certain senses and enhance performances within an ensemble. We collaborated on *Viewpoints AI* with a local experimental theater company, Out of Hand Theater, which employs Viewpoints as a system for exploring scene work and improvisation.
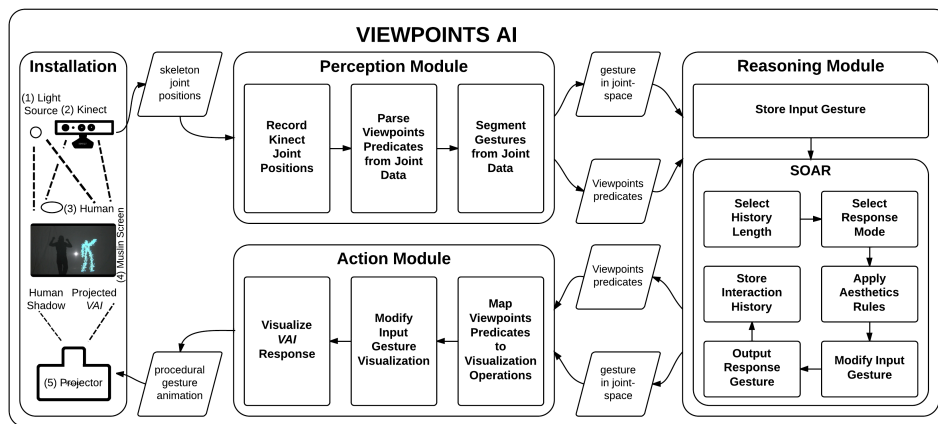
Figure 2: The interaction proceeds as follows: (1) The human's shadow is thrown by a light source, which illuminates the entire muslin screen. (2) The Kinect sensor records the human's gestures and is connected to the computer controlling the installation. (3) The literal performance of the human is not seen by spectators and is rendered as a shadow on the muslin screen. (4) The muslin sheet serves as a screen for the projection of VAI and as a semi-transparent layer for the shadow. (5) A projector connected to the control computer projects VAI onto the screen from the side of the spectators.

There are five Viewpoints categories (Bogart and Landau, 2005). *Space* perceptions are about seeing and feeling physical relationships; the actor can communicate using the empathy of the audience for *space*. The *shape* category concerns the view of physical forms shaped by bodies and the understanding of these forms in relation to other bodies or to the surroundings. *Time* perceptions are about time and its regulation. The fourth category is *emotion*, which is experiencing states of being and communicate them to the audience. Finally, the *movement* category concerns kinetic sensations as well as the manner of performing these motions: jerky/smooth, fast/slow and how the different parts of the body of the actor interact.

Out of Hand Theater uses Viewpoints technique to improve acting, improvisation skills, creativity and to learn to create original performances. Our collaborator at the Out of Hand Theater helped us to focus on using the Viewpoints of *time* and *space* in particular, for computational definition in the *Viewpoints AI* system. Within the Viewpoints of time, we chose to computationally define *tempo* (the rate at which movement occurs), duration (how long a movement or a sequence of movements continues before changing) and repetition (the amount of repetitions of a parameter for an internal or an external movement). We also introduced *energy* as additional parameter in the category *movement*, which seems to be an important perceptual characteristic of a gesture, consisting of the total amount of movement occurring.

In the near future, we plan to formalize more of the values in the Viewpoints of *space* category, however at present, we have used limited parameters from *shape* (combination of lines and curves outlining the bodies), *gesture* (behavioral or abstract movement made by a part or parts

of the body, i.e. a shape with a beginning, a middle and an end), *architecture* (physical surroundings of the performers and how awareness of it affects their movements), *spatial relationship* (distance between a body or a group of bodies and another body, another group of bodies or the architecture) and *topography* (landscape and its material defined by the movements of the performers).

The Viewpoints AI Installation

The *Viewpoints AI* installation is a juxtaposition of a virtual agent and human interactor onto the same liminal 2D plane. This is made possible by projecting onto a single canvas from both sides. The design of the *Viewpoints AI* installation allows for a gestural co-performance between a human and an AI agent called *VAI*. Spectators observe the installation from the front, viewing the human participant's shadow and *VAI*'s projection interacting together on screen.

A muslin screen is used here as a display, which is often used in shadow theater, since it provides a good quality projected image while being transparent enough to accurately see the shadows cast on it. The use of shadow play itself was designed to increase the presence of the experience for the human participant by being abstract enough to focus audience attention on both interactors while remaining analog enough for the human to produce a nuanced effect at light speed with the slightest movement. *Figure* 2 illustrates the installation and architecture of the *Viewpoints AI* system.

The *Viewpoints AI* installation is currently a turn-based system with alternating turns between the human and *VAI*, since creating fluid or naturalistic turn-taking in interactive

systems (with natural pauses, smooth interruptions and other social dynamics) is still an active research field (Chao and Thomaz, 2010) and working to incorporate state of the art turn-taking would have required changing the initial focus of the research. As a result, both the human and *VAI* take turns performing their gestures, where each turn is defined as the period of motion between two periods of stillness (thresholded period of time, based on current Kinect frame rate, where the amount of body movement detected is close to zero). The entire motion between the two periods of stillness is considered a single gesture. Future work will consider less discrete approaches to turns.

Computational Architecture

The *Viewpoints AI* system is composed of three conceptual parts, namely the *perception*, *reasoning* and *action* modules. The *perception* module senses the user's body and derives Viewpoints information from the user's gestural input. The *reasoning* module models the decision-making involved in choosing a response to a perceived gestural input. Finally, the *action* module visually represents *VAI*, the virtual interactor, and the gestural response.

Perception

At the front end of *Viewpoints AI* system, SimpleOpenNI, a wrapper for the OpenNI / NITE libraries, repeatedly extracts a simplistic representation of the human body from Kinect depth data. These simplistic representations consist of points in 3D space (called joints), including ones for head, shoulders, elbows, knees and so on. We define a gesture as the sequence of Kinect joint positions between two periods of stillness. These positions are used to calculate viewpoints predicates: predicates that indicate values of Viewpoints parameters according to our formalization of them for use in the reasoning module.

Viewpoints predicates are defined in terms of joint positions, speeds and accelerations of joint movements. In particular, current *tempo* is measured as maximum of joint speeds and discretized to produce *tempo* predicates (such as TEMPO_HIGH). *Energy* is defined through amount of movement, which is measured as sum of joint speeds; the difference between the two is that *energy* is the amount of movement while *tempo* is the rate at which that movement occurs. *Smoothness* is measured through maximum of joint accelerations, with high acceleration corresponding to low *smoothness* and vice versa. The values for calculating predicates are measured each several frames. For *tempo* and *energy*, these values are then averaged over a short time window, based on the current Kinect frame rate, to smoothen the values of predicates. The need for smoothing arises since both *tempo* and *energy* are perceived as belonging to rather prolonged intervals of time rather than

individual frames. For *smoothness*, we perform a similar procedure, but calculate maximum instead of average. The reason for this is that gesture is perceived as angular if accelerations are high even in some moments. After smoothed values for *tempo*, *energy* and *smoothness* are calculated, we average them over the whole gesture and discretize to compute *average tempo*, *average energy* and *average smoothness* predicates. These predicates are used by Soar (Laird 2012) to reason about gesture as a whole, while smoothed, instantaneous values of predicates are passed to visualization to control instantaneous visualization parameters such as the figure's color.

Reasoning

The reasoning module of *Viewpoints AI* reads the viewpoints predicates that represent Viewpoints information about the human performer's gestural input and reasons about how to respond to that input. The reasoning module consists of the Soar cognitive architecture and an external gesture library. The Soar cognitive architecture is used to reason about how to respond to the gesture utilizing the viewpoints framework, while the gesture library stores user gestural inputs to enable gesture matching.

Soar

The Soar cognitive architecture is a rule-based model of human cognition (Laird 2012). It contains components such as a working memory made up of working memory elements, a procedural long-term memory consisting of rules that fire based on the current working memory state and operators that modify the working memory elements. Soar continuously operates a decision cycle consisting of proposing operators to apply at a particular state and working memory context, elaborating knowledge pertaining to the current state and contents of working memory, selecting from the different operators using any elaborated preference knowledge and finally applying that operator to modify its internal or external state. The reasoning module uses Soar to respond to gestural inputs from the human participant based on the Viewpoints framework.

Each input gesture and response gesture are stored in an internal interaction history in order to serve as experience that *VAI* can draw from in future turns or performances. *VAI* first randomly selects the length of its interactional history to look at when deciding how to respond. This can be one of *zero history* (ignores the interaction history), *single history* (looks only at human's current input gesture) or *multiple history* (looks at interactional history). Based on the length of history selected *VAI* selects the response mode or how to respond to that input gesture. The response modes are *no operation*, *repeat input gesture*, *functionally transform input gesture*, *output new gesture* and *use response pattern*. The *no operation* response mode is used, when *VAI* should be still in its current position and not re-

act to the human gesture. The *repeat input gesture* response mode is used to mimic the human and copy their entire input gesture. The *functionally transform input gesture* response mode is used to modify the human's input gesture by transforming the viewpoints predicates associated with that gesture or by applying a functional transform on that gesture from a library of functional transforms. These functional transforms include transformations such as reflecting the direction of motion of a particular limb, switching the motions of two limbs and copying the motion of a limb to one or more remaining limbs. Transformations are done using the Viewpoints framework.

The *output new gesture* response mode draws from the internal interaction history in order to deliver novel (in the current interaction context) offers to the human as a base for equal co-creation of the gestural performance. In the current implementation of *VAI* this selection from the interaction history is conducted at random. However in future versions the selection is to be biased based on a measure of similarity between source and target viewpoints predicates such as similar *tempo* or duration.

The *use response pattern* response mode is currently under implementation and draws on pattern analysis and usage in order to decide how to respond. The current version is limited to utilizing the comedic rule of threes so that when dealing with repeated inputs of the same gesture from a human the first two times the system responds in the same way while breaking that expectation the third time.

Gesture Library

Viewpoints AI needed a dynamic system that stored a history of gestures already performed in some sort of a database, as well as accepts queries to quickly fetch past gestures in the form of joint space data to pass onto our visualization layer. This was required so as to notify *VAI* that a gesture from its past had been repeated by the human and to act accordingly either while using interaction patterns or while using aesthetics rules. The *Viewpoints AI* system's gesture recognition requirements were unique as well, since it was required to do fast matching over arbitrary gestures without training, albeit with assumptions made about starting positions for gestures. An effort to fulfill these requirements gave birth to the *Gesture Library*.

The *Gesture Library* is a hierarchical construct that dissects gesture frames into discrete *poses* (the set of instantaneous body joint positions in space with adjustable tolerances for unconscious body motion) and makes use of a *trie* data structure (Fredkin 1960) to populate a tree of gestures in order to provide quick gesture matching by traversing the *trie*. Gestures that share starting *poses* branch out from the same root nodes, wherein each node represents a single *pose* comprised of *Kinect* joint space data. Each time a new gesture is recognized, it is passed to the *Gesture Li-*

brary, which first slices the gesture from frames into *poses* and then performs a sequential tree search to find matching nodes (*poses*) in the existing library. When a partial match is detected, a fork is added to existing common nodes to generate a new gesture branch. If a gesture matches a previous one completely, no action is taken; and if there is no match at all, a completely new gesture branch is added to the tree's root. Each discrete gesture shares a unique ID that is used to quickly look up matching gesture data from a second data structure, the *Gesture Table*.

The *Gesture Table* exists in parallel to the *Gesture Library* and is relationally populated each time a new gesture is detected by the library. While the *Gesture Library* is lossy due to the representation of gestures as poses or neighborhoods of joint position sets, over time, it is optimal for matching functions via traversal in order to detect past gestures. In contrast, the *Gesture Table* provides a faster O(1) lookup using the unique ID while storing a full copy of the original gesture's joint space data.

Action: Visualization

VAI is presented as a cloud of bright moving particles (inspired by a flock of fireflies), dynamically forming the shape of human body (See Figure 1). The firefly cloud representation was selected as it matched a desired aesthetic that aims for a compromise between hyper-real and cartoonish, as well as between purely humanoid and completely ethereal. In addition, this representation made possible a range of visual effects useful for displaying values of viewpoints predicates.

The visualization component consists of four parts: a behavioral model of 'fireflies' (simple agents that turn towards 'sources of light'); gradients / fields of 'light' directions that attract the fireflies, constructed based on joints data; the *mimicry manager* that applies functional transformations to the performer's joints data; and a player, which feeds the *mimicry manager* with joint positions from the original gesture and allows *VAI* to tweak speed, duration, acceleration and repetition. All of these modules respond to viewpoints predicates and transformations received from Soar. Requests to perform functional transformations are processed by the *mimicry manager*. Relative *tempo* and duration transformations as well as repetition requests are handled by the player component. Viewpoint predicates are visualized by tweaking parameters of firefly and field models. In particular, *tempo* is shown through speed of firefly movements. Low *energy* of gesture is shown by red color of fireflies (inspired by glowing ashes), while high *energy* is shown by blue color (inspired by hot stars) and flashing of the fireflies (as if they were electrified by high-speed movement through air). High *smoothness* is represented by making the figure leave two types of trails. One is a glowing trail after each individual

firefly that is a representation of perceived trail left by a bright source of light in the dark. Fireflies lagging behind the moving figure create another trail. Elegant, long trails for user movements are perceived to be smoother.

Observations and Future Work

The prototype of the *Viewpoints AI* application was presented to a general audience during a research showcase. The performance setup was deployed in one of the labs opened to public access throughout the event. Incoming visitors were provided with explanations about the purpose of the setup and a short demonstration showing how to interact with the application. After that, they were invited to interact with the application by themselves. This format of the demo allowed us to collect valuable user feedback.

The overall reception of the *Viewpoints AI* system was positive. The concept and the visual aesthetics of the *Viewpoints AI* experience received highly appreciative comments. We observed several cases in which interaction was very engaging - people continued to play with *VAI* for long periods of time (sometimes more than 10 minutes) and stepped down only when forced to by other visitors, waiting for their turn. They challenged the system with complex gestures and then with interest observed *VAI*'s responses. There were also numerous cases where people interacted with the application in an equal and co-creative manner, by accepting the creative offers put forward by *VAI*. In one example, a little boy was gesturing at the system when it waved at him. The boy accepted the new offer and waved back in response. In another case, a participant was trying to grasp the look-and-feel of *VAI*'s movements and respond in the same way but with elaborated gestures. People were also reacting to "aggressive" gestures of the figure by shrinking back from it in a natural manner.

Critical remarks were of two types. First, people were not always able to capture the co-creative nature of the system. Some of them perceived *VAI* responses either as exactly copying human movements or as totally random and non-contextual (i.e. the Tale-Spin Effect (Wardrip-Fruin 2009)). This phenomenon highlights the challenge of having the AI's responses be both different enough and similar enough to the user's creative offer at the same time, which seems to be inherent for such co-creative applications. However, this challenge justifies the need for AI-based approaches in co-creative interactive performance. In people who were highly engaged in the interaction, we observed the opposite phenomenon - they attributed more meaning to the response than was actually present (i.e. the Eliza Effect (Wardrip-Fruin 2009)). One way to approach this challenge is to choose past gestures from *VAI*'s interaction history based on the similarity of its viewpoints predicates to the performer's current input gesture. This is analogous to the notion of jazz improvisers drawing riffs from long-term

memory according to semantic activation during real-time improvisation (Boden 2003).

The second aspect that caused audience dissatisfaction was the turn-based manner of interaction. People were frustrated by the need to wait for the response instead of seeing the reaction immediately. It seems that this frustration was caused by the disruption in communication between human and virtual character that usually happens in a turn-based interaction. While fluid human-computer turn-taking is an active area of research (Chao and Thomaz 2010), future versions of the *Viewpoints AI* system aim to provide a more fluid, naturalistic turn-taking system.

Finally, a promising path for future work is to do supervised learning of computational definitions for the Viewpoints framework for better scalability and accuracy. Gestures will be presented as points in multidimensional space of objective parameters (like average angular speeds of all joints). Classification algorithms will be trained on these parameters in order to identify spaces of gestures corresponding to different values of a viewpoints predicate. We hypothesize that similar techniques will also allow the system to learn some semantics of gestures - for instance, whether a gesture is aggressive, or whether it is a bow. Being able to infer such parameters would allow for a more meaningful interaction with *VAI*.

Contributions

Gestural interaction presents new opportunities for engaging experiences in interactive narrative and digital entertainment. However, the challenges associated with gestural interaction are numerous at present. The *Viewpoints AI* system is our approach to making gestural interaction more open-ended and expressive. Procedural gesture interpretation could improve the scalability of creating a gestural interface by reducing the authoring bottleneck for instantial gesture assets. A procedural approach to gesture interpretation such as the use of Viewpoints could thus result in more expressive player control in games and a more natural interaction modality for interactive digital experiences.

Acknowledgements

References

Boden, Margaret A. *The creative mind: Myths and mechanisms*. Routledge, 2003.

Bogart, A., and T. Landau. 2006. *The Viewpoints Book: A Practical Guide to Viewpoints and Composition*. Theatre Communications Group. http://books.google.com/books?id=7WJQAAAAMAAJ.

Calvert, Thomas W., Christopher Welman, Severin Gaudet, Thecla Schiphorst, and Catherine Lee. 1991. "Composition of Multiple Figure Sequences for Dance and Animation." *The Visual Computer* 7 (2-3): 114–121.

Chao, Crystal, and Andrea L. Thomaz. 2010. "Turn Taking for Human-robot Interaction." In *AAAI Fall Symposium on Dialog with Robots*, 132–134.

Chen, Fan Pen Li. 2007. *Chinese Shadow Theatre: History, Popular Religion, and Women Warriors*. McGill-Queen's Press - MQUP.

Christoulakis, Marios, Andreas Pitsiladis, Argiro Moraiti, Nektarios Moumoutzis, and Stavros Christodoulakis. "EShadow: A Tool for Digital Storytelling Based on Traditional Greek Shadow Theatre."

Colton, Simon. 2008. "Creativity Versus the Perception of Creativity in Computational Systems." In *Proceedings of the AAAI Spring Symposium on Creative Systems*.

Cowell, S., K. Warwick, and J. Raff. "America's got talent [Television series]." *London: FremantleMedia Licensing Worldwide (2006)*.

El-Nasr, Magy Seif. 2007. "Interaction, Narrative, and Drama: Creating an Adaptive Interactive Narrative Using Performance Arts Theories." *Interaction Studies* 8 (2): 209–240.

Fredkin, Edward. 1960. "Trie Memory." *Communications of the ACM* 3 (9): 490–499.

Hecker, Chris, Bernd Raabe, Ryan W. Enslow, John DeWeese, Jordan Maynard, and Kees van Prooijen. 2008. "Real-time Motion Retargeting to Highly Varied User-created Morphologies." In *ACM Transactions on Graphics (TOG)*, 27:27. ACM.

Hodhod, Rania, Andreya Piplica, and Brian Magerko. 2012. "A Formal Architecture of Shared Mental Models for Computational Improvisational Agents." In *Intelligent Virtual Agents*, 440–446. Springer.

Laban, Rudolf, Roderyk Lange, and Diana Baddeley. *Laban's principles of dance and movement notation*. Macdonald and Evans, 1975.

Laird, John E. 2012. *The Soar Cognitive Architecture*. MIT Press.

Latulipe, Celine, David Wilson, Sybil Huskey, Berto Gonzalez, and Melissa Word. 2011. "Temporal Integration of Interactive Technology in Dance: Creative Process Impacts." In *Proceedings of the 8th ACM Conference on Creativity and Cognition*, 107–116. ACM.

Magerko, Brian, Waleed Manzoul, Mark Riedl, Allan Baumer, Daniel Fuller, Kurt Luther, and Celia Pearce. 2009. "An Empirical Study of Cognition and Theatrical Improvisation." In *Proceedings of the Seventh ACM Conference on Creativity and Cognition*, 117–126. ACM.

Mateas, Michael, and Andrew Stern. 2003. "Façade: An Experiment in Building a Fully-realized Interactive Drama." In *Game Developers Conference, Game Design Track*, 2:82.

McCorduck, Pamela. 1991. *Aaron's Code*. WH Freeman & Co.

O'Neill, Brian, Andreya Piplica, Daniel Fuller, and Brian Magerko. 2011. "A Knowledge-based Framework for the Collaborative Improvisation of Scene Introductions." In *Interactive Storytelling*, 85–96. Springer.

OpenNI. http://www.openni.org

Overlie, Mary. 2006. "The Six Viewpoints." *Training of the American Actor*: 187–221.

Perlin, Ken, and Athomas Goldberg. "Improv: A system for scripting interactive actors in virtual worlds." In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pp. 205-216. ACM, 1996.

Pinhanez, Claudio S., and Aaron F. Bobick. ""It/I": a theater play featuring an autonomous computer character." *Presence: Teleoperators and Virtual Environments* 11, no. 5 (2002): 536-548.

Piplica, Andreya, Christopher DeLeon, and Brian Magerko. 2012. "Full-body Gesture Interaction with Improvisational Narrative Agents." In *Intelligent Virtual Agents*, 514–516. Springer.

PrimeSense. NITE Middleware. http://www.primesense.com/?p=515

Riedl, Mark O., and Vadim Bulitko. "Interactive Narrative: An Intelligent Systems Approach." *AI Magazine* 34, no. 1 (2013): 67-77.

Sharma, Megha, Dale Hildebrandt, Gem Newman, James E. Young, and Rasit Eskicioglu. "Communicating affect via flight path Exploring use of the Laban Effort System for designing affective locomotion paths." In *Human-Robot Interaction (HRI), 2013 8th ACM/IEEE International Conference on*, pp. 293-300. IEEE, 2013.

Subyen, Pattarawut, Diego Maranan, Thecla Schiphorst, Philippe Pasquier, and Lyn Bartram. 2011. "EMVIZ: The Poetics of Movement Quality Visualization." In *Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, 121–128. ACM.

Wardrip-Fruin, Noah. *Expressive Processing: Digital fictions, computer games, and software studies*. The MIT Press, 2009.