A Knowledge-Based Framework for the Collaborative Improvisation of Scene Introductions

Brian O'Neill[1], Andreya Piplica[1], Daniel Fuller[2], Brian Magerko[2]

[1] School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA
[2] School of Literature, Communication, and Culture, Georgia Institute of Technology, Atlanta, GA, USA
boneill@cc.gatech.edu, apiplica3@gatech.edu, gentristar@gmail.com, magerko@gatech.edu

Abstract. This article describes a framework for the mixed-initiative collaborative creation of introductions to improvised theatrical scenes. This framework is based on the empirical study of experienced improvisational actors and the processes they use to reach shared understanding while creating the scene. Improvisation is a notable creative act, where the process of creating the scene is as much a product as the scene itself. Our framework models the processes of narrative scene establishment. It is designed to allow for the collaborative co-creation of the narrative by both human and computational improvisers. This mixed-initiative approach allows either type of improviser (AI or human) to deal with the ambiguities that are inherent to improvisational theatre. This emphasis on equal collaborative creation also differentiates this framework from existing work in story generation and interactive narrative.

Keywords: Collaborative environments for interactive storytelling, semantic knowledge for interactive storytelling, virtual characters and agents

1 Introduction

Improvisation is a well-known form of several types of entertainment, including music and theatre. Because of its ubiquity in creative domains, improvisation is a growing area of interest for researchers in the computational creativity field. Improvisation (or "improv") in an artistic domain (e.g. improvisational theatre) has been defined as the creation of an artifact and/or performance with aesthetic goals in real-time that is not completely prescribed in terms of functional and/or content constraints [1]. Studying improvisation, and improvisational theatre, in particular, provides a unique perspective on human creative processes and narrative, informing the development of computational agents capable of improvisation.

Our previous work on the study of improvisational actors has culminated in the construction of an agent-based framework for playing *Party Quirks*, a well-known improvisational game, in collaboration with human improvisers [2]. *Party Quirks* was selected as the initial domain for developing an improvisational framework in part because the game inherently lacks a narrative aspect; it focuses on the process of cognitive convergence (improvisers "getting on the same page") without concern for telling a coherent story. The work presented in this paper builds on our previous work by introducing a knowledge-based framework for collaboratively creating the elements present in the introduction of an improvised scene with a narrative (i.e. the *platform* of a scene). The platform defines four aspects of a scene: the location of the scene, the characters within the scene, the relationship between those characters, and the shared (joint) activity which the characters are participating in.

Generating an improvised scene differs significantly from canonical approaches to story generation. First and foremost, the processes involved in improvising the story are appreciated as much as, if not more than, the story itself. The entertainment value of improv comes from watching the formation of a narrative directly in front of the audience and not solely from viewing a performed scene (as in traditional theatre). In other words, *the process is part of the product* in improvisation. Second, improv performances contain inherent ambiguities that must be resolved by both the audience and the actors. Finally, improvisation is an imperfectly coordinated multi-agent process, whereas canonical story generation is single-agent. Each improviser has a different conception (or *mental model*) of the narrative and can only share this model through the performance. This lack of coordination between the multiple authors of the scene is an additional source of ambiguity in improvisational performances. This differs from less improvised settings (e.g. a screenwriting meeting) where coordination is less constrained. In generating the platform, we need to understand how improvisers reach a shared model of the scene through performance alone.

The work presented in this paper is part of ongoing research into improvisational actors, focusing on gaining an understanding of human creativity with respect to the collaborative creation of stories. We then apply this understanding to developing intelligent improvisational agents. As part of this effort, we studied performances by experienced improvisational actors. Based on data collected from these performances, we introduce a new conceptual framework for the co-creation of the platform of a scene by human and computational improvisers. This joint co-construction occurs within a game of *Three Line Scene*, an improvisational teaching game used for learning how to establish the platform quickly. In this game, two actors have to create a complete platform for an improvised scene in only three lines of dialogue. The framework defines the agent's knowledge representation as well as the processes the agent applies to interpret ambiguous information and map those interpretations to instantiated facts about the reality continuously being co-constructed on stage.

2 Related Work

We have studied how improvisers deal with *cognitive divergences* (i.e. when actors are not "on the same page") within the context of the improv game *Party Quirks* [3]. Although *Party Quirks* contains no narrative development, it is a game that clearly illustrates the types of offers (presentations made by improvisers) that actors use to resolve divergences on stage. The same processes can be used to resolve divergences in narrative-oriented scenes as well [1].

Our current work uses the game *Three Line Scene* to apply our studies of human improvisers to a narrative context. In this game, two actors take turns presenting dialogue and motions with the goal of establishing a complete platform in three turns. This game forces actors to make strong offers that contribute multiple elements to the platform at once and to accept and augment previous presentations [4]. In addition to dialogue offers, actors utilize distinct motions, which can convey information about a character or the joint activity [5]. We use *Three Line Scene* as the basis for our framework due to its simple rules, generative capabilities, and focus on the platform.

Most improvisation research has focused on music [6, 7]. Emergent music creation work [8, 9] has led to improvisational agents that can co-create in a musical performance [10]. Theatrical improvisers cannot rely on explicit meta-communication and rarely use pre-established structures (i.e. stock characters, narrative structure, etc.) analogous to those available to improvisational musicians (i.e. chord progressions, key signatures, etc.). Sociolinguistic studies of theatrical improvisation have found that all of a theatrical improviser's actions and dialogue are generated and presented within the performance as offers for the scene [11]. The collective responses to these offers – which can be accepted or rejected, augmented or redirected [4] – create the improvised narrative. Implementations of theatrical improvisation agents [12-15] have typically focused on portraying particular aspects of improvisation informed by classic improvisation texts rather than on creating narratives. The fields of story generation and interactive narrative can benefit from applications of improvisational techniques to create interesting narratives without predetermined planning [16].

Story generation employs intelligent agents to create stories. Story generation systems typically use a single agent [17] or multiple agents capable of communicating about the content and the presentation of the narrative [18]. In contrast, our work focuses on multiple agents creating a story without explicit communication within the context of a performance.

Interactive narratives incorporate humans into the story process by allowing users to influence the path and outcome of an adaptable story. Interactive narrative systems typically have fixed, pre-authored story elements that make up the atomic elements of co-creating a story. Human interactors influence the selection of atoms through various mechanisms (e.g. navigating the social space of a story world [19] or uttering dialogue that maps to positions on the story's Aristotelian arc [20]). In some cases, agents in interactive narrative systems have advance knowledge of the characters and an intended narrative. In other cases, agents do not contribute directly to the creation of the emergent narrative but influence the outcome from a directorial role [21]. A common factor among these systems is that although humans and agents co-create a story, they do not do so as equals. Our framework lays the foundation for agents that equally co-create an emergent narrative with a human interactor, where each player adds something new to the narrative with each play they make.

3 Platform Creation Framework

We have developed a framework for implementing an agent capable of playing a modified version of *Three Line Scene*, which omits the aspects of narrative that occur later in scenes. We note that this framework is a human behavior model, rather than a cognitive model, of improvisation. That is, we do not claim that human improvisers carry out the exact processes described below. We do, however, believe that this framework models the knowledge necessary to establish the platform of a scene. The omission of narrative aspects beyond the platform allows us to focus on the collaborative construction of the platform as an initial point in developing agents that can co-construct narratives with humans or other agents. The computational agent uses a modified version of Sawyer's definition of "platform" [11] to reason about the state of the emerging scene. Sawyer's definition, derived from the common practices of improvisers, identifies four elements of the platform: the *characters* in the scene, their *relationship* to each other, the *joint activity* they are engaged in, and their *location*. Sawyer notes that relationship and location can often be inferred from the characters and joint activity, respectively. Therefore, our agent's knowledge structure focuses on reasoning about the characters and joint activity. This reduces both the amount of authoring and computation needed to reason about the platform.

In addition to the character and joint activity aspects of the platform, our knowledge structure contains other types of elements that are necessary for interpreting the scene and making presentations. The first of these are *motions* and *actions*. A motion is a physical representation of a movement while an action is an intended or interpreted meaning of a motion. For example, an agent may hold out one hand shaped as a loosely closed vertical fist. In this case, the motion is *hold out fist*, while the action may be *give bag*. A single motion can portray multiple actions, so holding out a fist in this manner could also be interpreted as *give bottle*. Finally, our knowledge structure contains icons that convey aspects of a scene that would typically be presented through dialogue rather than through some physical motion. Such aspects include presentations about another improviser's character and explicit instantiation or clarification of an element in an improviser's mental model.

The agent's knowledge base is divided into four *categories*: motions, actions, characters, and joint activity. The agent also organizes its knowledge of the scene based on which improviser it is associated with – that is, it keeps knowledge of its own character, actions, and motions distinct from those of the other improviser. Each computational agent tracks the *current instantiation* (motions or icons that either improviser has presented) and its own *mental model* (elements it supposes are in the scene that have not yet been explicitly confirmed). The agent treats elements in its mental model as true until provided with evidence (a presentation from the other improviser) to the contrary. The agent uses the elements it supposes to be true when expanding its mental model. It does not consider other elements in knowledge categories where it already supposes some element to be true. Actions that the agent chooses to present are a special case; the agent will ignore actions already in its mental model so as to avoid repeating actions. The agent will always consider instantiated elements to be true.

Because the categories of our knowledge structure are highly inter-related, we need an approach to show approximately how related any two items are. We have adopted a fuzzy logic approach similar to the one used in the *Party Quirks* framework [2] to represent the association between two elements in our knowledge structure. Every element within our knowledge structure is connected to elements in related categories with a degree of association (DOA), which could be considered a bi-directional version of degree of membership (DOM) in fuzzy logic. Whereas DOM is unidirectional, where some element is a member of a set to some degree, DOA is a bi-directional relationship representing the extent to which two items are related, where 0 means that the items have no association whatsoever, and 1 means that the items are highly correlated with each other. Table 1 shows a sample DOA table of the degrees of association between possible characters and joint activities.

The agent assumes that the other actor intends to communicate their mental model clearly in order to reach a shared understanding about the scene [3]. Therefore, when expanding its mental model, the agent adds one of the most iconic interpretations of what it has seen or what its mental model implies. Iconicity (or its inverse, ambiguity) is a measure of the uniqueness of the DOA values between a single element and every element of another category. (See [2] for greater detail on how ambiguity is calculated.) For example, suppose the agent is considering what it knows about the other actor's character given that it thinks the joint activity is *gambling*. If most characters have a medium degree of association with gambling, then these characters are not iconic gamblers. In this case, a character with a very high or very low degree of association with gambling would be an iconic, and therefore preferable, choice. The agent only considers elements that are among the most iconic possible interpretations. Considering iconicity first ensures that the agent presents a choice that is easily distinguished by the other agent regardless of its DOA with other elements in its mental model.

3.1 Simplifications

We have made some simplifying assumptions for the purposes of modeling a computational improviser. First, we made the decision that the agent's mental model does not include theory of mind – that is, the agent does not track what it supposes is in the other improviser's mental model of the scene. Agents can use theory of mind to model other interactors, evaluate the outcome of actions, and update goals [22]. In *Three Line Scene*, models of other interactors would provide insight into the causes of and possible resolutions to cognitive divergences. We find it valuable to evaluate whether our agent can correct divergences without keeping a model of the other actor's behavior before attempting to implement a theory of mind system.

Furthermore, agents are limited to communicating with motions and icons, where one icon represents one element of the agent's mental model. Icons avoid the issue of addressing natural language generation and processing. This allows us to focus on the cognitive process of setting up a platform rather than the detailed mechanics of verbal communication. Using icons avoids the use of "canned" language, while still allowing

Table 1. Sample knowledge structure table for the *Tiny West* domain shows the degrees of association that different characters have with different joint activities.

	Robbery	Showdown	Drinking	Apprehending	…
Outlaw	0.9	0.8	0.8	0.8	…
Gunslinger	0.7	0.9	0.8	0.6	…
Sheriff	0.8	0.8	0.7	0.9	…
Banker	0.8	0.2	0.5	0.4	…
…	…	…	…	…	…

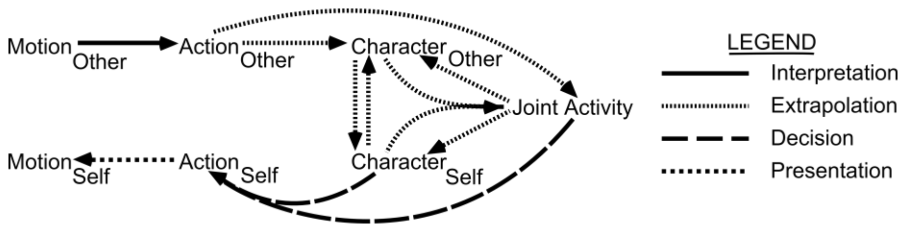

Figure 1. Organization of knowledge structure categories throughout the turn process.

for the direct communication of ideas between two agents (i.e. communication that does not need to be interpreted through ambiguous motions and actions). While the framework could accommodate ambiguous icons (i.e., icons that might represent multiple elements), this would greatly increase the complexity of authoring. In our observations of improvisers, ambiguous motions tended to create more divergences in setting up the platform than ambiguity of language, so we do not feel that this added complexity will be valuable at this stage. In addition, the use of icons also simplifies the interactions between human improvisers and the computational agent. It provides a common set of symbols that both the human interactor and the computational agent can understand. Ideally, the icons will be designed so that the human can understand what they represent and be able to use them with little to no instruction.

3.2 Turn Process

This section describes the process our narrative agents take in our modified version of *Three Line Scene*. We presume that agents use this process throughout the game, except during the first turn. In our analysis of improv actors in unconstrained narrative scenes, we noticed that actors would start scenes with potentially ambiguous motions rather than dialogue. Additionally, this turn process is based on interpreting previous offers and extrapolating from knowledge about the scene, neither of which is possible when considering the first move of the scene. Hence, we presume that an improviser selects an arbitrary motion and presents it as the initial move in the performance.

A turn in this framework is divided into five phases: *perception*, *interpretation*, *extrapolation*, *decision*, and *presentation*. The agent *perceives* the previous presentation, *interprets* the motions and/or icons and checks for divergences, *extrapolates* the interpretations to build its mental model of the scene, *decides* which mental model elements to present, and *presents* the elements with a motion, an icon, or both. Figure 1 summarizes our knowledge structure and shows which element categories can be reached from other categories during each turn phase. The details of each phase are discussed below.

In order to visualize the turn process, it is helpful to focus on a limited data set. For such purposes, we present *Tiny West*, an example story world designed for the purpose of illustrating our framework. *Tiny West r*epresents a limited set of characters, actions, and joint activities typical of the western genre in cinema. The iconic nature of the elements of this genre makes it valuable for highlighting strong and weak degrees of association. Furthermore, given the cultural prevalence of westerns, the relationships between the elements are fairly well-known and consistent across media, leading to less debate over specific values for the sake of authoring.

Perception. The agent receives as input whatever was presented by the other agent in the previous turn. We presume that our agents have perfect perception; that is, they observe icon and motion presentations without error. This perfect perception eliminates the uncertainty that arises in viewing a presentation so that agents can focus instead on the semantic ambiguity of a presentation. Motions remain open to interpretation about which action they represent. Icons represent a single element of the knowledge base, so they can be directly added to an agent's mental model.

In our *Tiny West* example, George opens the scene by randomly choosing a motion, presenting the motion *quickdraw*. As explained above, choosing a random motion to begin with is a legitimate move for theatrical improvisers. Ann perceives the *quickdraw* motion and adds it to her mental model.

Interpretation. Once a presentation is perceived, an agent must interpret its meaning. A single motion can potentially portray multiple actions, so the agent must select a single interpretation for that motion. First, the agent finds all possible actions that the motion may have been intended to portray. As described in Section 3, the agent only considers iconic elements when expanding its mental model. This ensures that the agent selects something that is likely to be what the other actor intended. The agent probabilistically selects from the iconic interpretations, favoring those with higher DOA. The agent adds the interpretation to its mental model as part of what the other actor has contributed to the scene.

In the *Tiny West* example, Ann interprets the *quickdraw* motion. She considers iconic actions with given the *quickdraw* motion. The iconic actions are *draw gun* (DOA 0.8), *show off* (DOA 0.6), and *stick up* (DOA 0.9). Ann probabilistically chooses between the three. She interprets *stick up* as the intended action and adds it to her mental model.

When an agent interprets an icon, the icon may introduce a divergence in the agent's mental model. Such a divergence occurs when the instantiated element conflicts with some other element in the same category in the agent's mental model. In this case, the instantiated element supersedes the mental model element. Once the agent replaces an element of its mental model, it must update the rest of the model to ensure that all elements are still associated with each other, as the newly instantiated element may have a DOA of 0 with something already in the mental model. The agent sets an acceptability threshold when it considers the association between elements of two categories. For all categories where the agent has an element in its mental model, the association between elements in connected categories must be above the threshold. If they are not, the agent replaces elements below the threshold with new elements that are above it, selecting probabilistically based on DOA.

Suppose in a later turn that George has *showdown* as the joint activity in his mental model, but this has not been instantiated yet. If Ann presents an icon for the joint activity *robbery*, George removes *showdown* from his mental model and replaces it with *robbery*. This resolves the divergence [3], and George can move on to extrapolation.

Extrapolation. The agent considers what else may be true in the scene based on the updated contents of its mental model in order to add new information to the scene [4]. This phase is not meant to simulate the exact cognitive processes that a human improviser uses; rather, it approximates the behavior of making inferences from previous presentations. The agent arbitrarily selects an interpreted element and tries to extrapolate to an element in another knowledge category. The choice of where to start from is ultimately insignificant, as every element of the current instantiation and the agent's mental model feeds into the extrapolation process. Given the selected element, the agent selects one of the two reachable categories to consider after accounting for what has been instantiated in the scene (e.g. if roles for both characters have been instantiated, then the agent will not consider characters). The selection is arbitrary when both categories are available – that is, when nothing from either category has been instantiated yet. In practice, as the scene continues and more platform elements are instantiated, it becomes increasingly likely that one of the categories cannot be selected. From the available category, the agent selects an element based on its iconicity and DOA as described in Section 3.

Extrapolations influenced by multiple categories need to be represented compactly to reduce authoring. The agent takes the fuzzy AND (i.e. a minimum function) [23] of the DOA between its own character and a potential joint activity and the DOA between the other improviser's character and that joint activity. This prevents the need for authoring an association cube between both sets of characters and joint activities. If the agent has not added both characters to its mental model, it does a regular extrapolation from whichever character it has in its mental model. Similarly, the other improviser's motions and actions both inform what character they may be, since different characters portray the same action with different motions. To represent this without an association cube between motions, actions, and characters, the agent extrapolates from an action to a character with a DOA approximately the same (within one standard deviation) as the DOA between that action and the initial motion.

We presume that the agent does not necessarily try to define every aspect of the platform during each turn. To do so, the agent would make several assumptions based on elements that are only in its mental model. The fewer assumptions the agent makes, the less likely it is to encounter divergences later. After each extrapolation, the agent decides whether to continue extrapolating. This decision is weighted based on how many extrapolations the agent has made on this turn. Continuing to extrapolate becomes less likely with each pass. If the agent makes another pass at extrapolation, it returns to the beginning of the process, working this time from the newly extrapolated element. If the agent ends the extrapolation phase, it moves to the *decision phase*.

Continuing with our *Tiny West* example, Ann extrapolates from George's action. Actions extrapolate to joint activities (see Figure 1), so Ann considers which joint activities are iconic given *stick up*. Both *robbery* and *showdown* are iconic given *stick up*; these activities have high DOA (0.9 and 0.7, respectively) while all others have low DOA. Ann probabilistically selects *robbery*. She decides to continue extrapolating. From joint activities, she can extrapolate to either her own character or George's character. She decides to extrapolate to her character. *Outlaw* and *banker* both have high DOAs with the joint activity *robbery* (0.9 and 0.8, respectively), while other characters tend to have mid-range DOAs. Both are iconic. Ann probabilistically selects *banker*. Since Ann has now extrapolated twice, she is less likely to continue extrapolating. She decides to stop extrapolating now.

Decision. The agent decides to present an icon (a substitute for dialogue), an action, or some combination of the two. If the agent only has the other improviser's character in its mental model, it must present that with an icon. (An action cannot convey information about the other improviser's character.) Otherwise, there is no preference among the three choices, as each adds to the scene and would be valid for a human improviser. The agent arbitrarily decides which of these to present. If it chooses to present an icon, the agent selects an icon representing an arbitrarily chosen element of its mental model. If it chooses to present an action, the agent takes the fuzzy AND of the DOA between its own character and an action and the DOA between the joint activity and that action. Potential actions must be iconic for both the agent's character and the joint activity. If the agent does not have both its own character and the joint activity in its mental model, it selects an action based on the elements that are in its mental model. Finally, the agent adds the selected action to its mental model.

In *Tiny West*, Ann must now decide which aspects of her mental model to present and how. She decides to only present an action. She considers the actions that both her character *banker* and the joint activity *robbery* are associated with. The relevant actions are *give money* and *stick up*. Both *banker* and *robbery* are iconic and have high DOA with *give money* (both DOA 0.8, while other characters have low DOAs). The fuzzy AND of their DOAs with *give money* is 0.8. *Banker* is moderately associated with *stick up* (0.6), but so are most other characters, so this DOA is not iconic. *Robbery* is iconic and has high DOA with *stick up* (0.9 while most other activities have a low DOA). However, since the combination of *banker* and *stick up* is not iconic, Ann does not consider the *stick up* action further. Ann selects *give money*, as it is the only viable action for her to choose, and adds it to her mental model.

Presentation. If the agent decided to present an icon, it displays the icon it has selected. If the agent decided to present an action, it first converts that action into a motion. As mentioned earlier, different characters portray the same action with different motions. Thus both the agent's character and the selected action affect the motion it presents. Like in the extrapolation phase, the agent presents a motion that has approximately the same degree of association (within one standard deviation) with the action as that action has with the agent's character. (If the agent's character is not part of its mental model, the agent selects a motion as if it were extrapolating from actions to motions.) Again, this avoids the need for a motion-action-character association cube. The agent displays its motion, adds the icon and/or motion to its mental model, and concludes its turn.

Having chosen an action to present (*give money*), Ann must choose which motion to use to present that action. *Give money* is strongly associated with *hand over bag* (0.6) and *hand over money* (0.9). The DOA from *give money* to *hand over money* is closer to the DOA from *banker* to *give money* (0.8) than the DOA from *give money* to *hand over bag*. Ann chooses *hand over money*, which she then presents to George.

4 Discussion

This framework models the improvisation process for the construction of the platform in a narrative scene. In improvisational theatre, the process used by improvisers to create a scene is at least as important as the resulting scene itself. Our work is an extension of improvisational agents built for non-narrative games, based on the analysis of experienced improvisational actors [1]. While this work does not describe a cognitive model of improvisers, it does describe the behaviors used to establish the platform in a narrative scene.

Our framework represents the collaborative creation of the platform by both human and computational improvisers. This equal partnership in co-creating a story improves upon story generation and interactive narrative literature. We are presently unaware of any mixed-initiative story generation systems, although work has been proposed for mixed-initiative control of believable agents in interactive digital storytelling [24, 25]. Interactive narrative systems rely on pre-authored content. While users can affect the direction that an interactive narrative takes, they remain beholden to the content that has been written. In our approach, human users are true co-creators; there is no existing narrative content for the user to follow.

This framework currently assumes that improvisers, both computational and human, make only "good" plays. This is a reasonable assumption, as human improvisers are trained to advance the scene with each decision they make. In its most basic form, advancing the scene consists of accepting the previous offer and adding something else to the scene. Given this aspect of improvisation training, we designed our framework so that a computational improviser only makes good plays. While human interactors can still make "bad" plays, a computational agent can deal with an unusual offer from a human player by treating the offer as a divergence and adapting its play accordingly, as described in the Interpretation and Extrapolation phases.

A computational agent's ability to improvise inherently relies on its knowledge base. One might argue that a human improvising with a computational agent using this framework is just as limited by authored content as someone taking part in any other interactive narrative system. However, in an interactive narrative, the agent works from a set of preconceived story elements. In this framework, a computational agent works from a known knowledge base that is not connected to any portion of a narrative. While we do intend to limit interaction so that a human improviser cannot easily stray from the knowledge base, we still believe that a human improviser/co-creator is significantly less constrained than they would be in an interactive narrative.

We have described *Tiny West*, an example story world, to illustrate the framework and the agent turn process. *Tiny West* contains four characters, five joint activities, seven actions, and 14 motions, which means authoring approximately 200 DOA ratings. The DOA authoring requirements for this system are exponentially related to the number of elements in the knowledge base. However, these concerns do not affect the processes we have described. The improvisation turn process scales regardless of the size of the knowledge base.

We plan to evaluate this framework by asking people to improvise with computational agents. We used a similar evaluation technique with our *Party Quirks* system, including evaluation by a panel of experts at the 2011 Chicago Improv Festival [2]. People without improvisation experience will provide feedback about the interaction experience, particularly the use of icons. Participants will be asked if, at the end of the improvised scene, they could identify the elements of the platform that had been established. If participants describe the actual elements of the platform, then these results will validate the framework. We also plan to ask expert improvisers to compare the actions of the computational agents to the decisions they would have made in the same situation. Similarly, we intend to ask whether the framework prevented them from taking an action that they wanted to take (excluding authoring limitations). Ultimately, however, this framework is not an end. This work is one piece in understanding human creative processes with respect to improvisational theatre. Future work in conceptual blending, declarative knowledge, and procedural tacit knowledge for improvisation will all add to the richness of this work.

Improvisational theatre is a unique source of information for the study of human creative and cognitive processes. By studying expert improvisers, we have gained a greater understanding for how humans collaboratively create a narrative and how ambiguities in the creation are identified and repaired. The development of this framework is one step in a continuing research effort to create intelligent agents capable of performing these same processes in concert with humans. Future work will further explore narrative, investigating how an established platform affords the collaborative mixed-initiative creation of plot. With such an understanding, we will be able to develop agents that can improvise a complete narrative-rich scene.

References

1. Magerko, B., Manzoul, W., Riedl, M., Baumer, A., Fuller, D., Luther, K., Pearce, C.: An Empirical Study of Cognition and Theatrical Improvisation. In; Proc. Seventh ACM Conference on Creativity and Cognition. pp. 117–126. ACM, New York (2009).
2. Magerko, B., Dohogne, P., DeLeon, C.: Employing Fuzzy Concepts for Digital Improvisation Theatre. In: Proc. Seventh Annual Conference on AI and Interactive Digital Entertainment (AIIDE). AAAI Press (2011).
3. Fuller, D., Magerko, B.: Shared Mental Models in Improvisational Theatre. In: Proc. Eighth ACM Conference on Creativity and Cognition. ACM, New York (2011) (in press).
4. Baumer, A., Magerko, B.: An analysis of narrative moves in improvisational theatre. In: Aylett, R., Lim, M., Louchart, S., Petta, P., Riedl, M. (eds.) Interactive Storytelling. LNCS, vol. 6432, pp. 165–175. Springer, Heidelberg (2010).
5. Johnstone, K.: Impro for Storytellers. Routledge, New York (1999).
6. Reinholdsson, P.: Making Music Together: An Interactionist Perspective on Small-Group Performance in Jazz. Acta Universitatis Upsaliensis: Studia Musicologica Upsaliensia, Nova Series 14, Uppsala University (1998).
7. Alterhaug, B.: Improvisation on a triple theme: Creativity, Jazz Improvisation and Communication. Studia Musicologica Norvegica. 30, 97-118 (2004).
8. Beyls, P.: Interaction and Self-organisation in a Society of Musical Agents. In: Proc. ECAL 2007 Workshop on Music and Artificial Life. Lisbon, Portugal (2007).
9. McCormack, J.: Eden: An Evolutionary Sonic Ecosystem. In: Kelemen, J., Sosík, P. (eds.) Advances in Artificial Life, LNCS, vol. 2159, pp. 133–142. Springer, Heidelberg (2001).
10. Eigenfeldt, A.: Coming together: Composition by negotiation. In: Proc. International Conference on Multimedia. p. 1433–1436. ACM, New York, NY, USA (2010).
11. Sawyer, R.K.: Improvised dialogues: Emergence and creativity in conversation. Ablex Publishing Corporation, Westport, CT (2003).

12. Aylett, R.S., Louchart, S., Dias, J., Paiva, A., Vala, M.: Fearnot!: An experiment in emergent narrative. In: Panayiotopoulos, T., Gratch, J., Aylett, R., Ballin, D., Olivier, P., Rist, T. (eds.) Intelligent Virtual Agents. LNCS, vol. 3661, pp. 305–316. Springer, Heidelberg (2005).
13. Riedl, M.O., Stern, A.: Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) Technologies for Interactive Digital Storytelling and Entertainment. LNCS, vol. 4326, pp. 1--12. Springer, Heidelberg (2006).
14. Perlin, K., Goldberg, A.: Improv: A System for Scripting Interactive Actors in Virtual Worlds. In: Proc. 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96). pp. 205-216. ACM, New York (1996).
15. Bruce, A., Knight, J., Listopad, S., Magerko, B., Nourbakhsh, I.R.: Robot Improv: Using Drama to Create Believable Agents. In: Proc. IEEE International Conference on Robotics and Automation, 2000. pp. 4002 - 4008. IEEE, San Francisco, CA (2000).
16. Swartjes, I., Vromen, J.: Emergent Story Generation: Lessons from Improvisational Theater. In: Proc. AAAI Fall Symposium on Intelligent Narrative Technologies, Arlington, VA (2007).
17. Pérez y Pérez, R., Sharples, M.: Three computer-based models of storytelling: BRUTUS, MINSTREL and MEXICA. Knowledge-Based Systems. 17(1), 15-29 (2004).
18. Theune, M., Faas, S., Faas, E., Nijholt, A., Heylen, D.: The Virtual Storyteller: Story Creation by Intelligent Agents. In: Proc. Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference. pp. 204--215 (2002).
19. McCoy, J., Treanor, M., Samuel, B., Tearse, B., Mateas, M., Wardrip-Fruin, N.: Comme il Faut 2: A fully realized model for socially-oriented gameplay. In: Proc. Intelligent Narrative Technologies III Workshop. pp. 1-8. ACM, New York, NY, USA (2010).
20. Mateas, M., Stern, A.: Integrating plot, character and natural language processing in the interactive drama Façade. In: Proc. 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment. Darmstadt, Germany (2003).
21. Zhang, L., Gillies, M., Barnden, J.A., Hendley, R.J., Lee, M.G., Wallington, A.M.: Affect detection and an automated improvisational AI actor in E-drama. In: Huang, T., Nijholt, A., Pantic, M., Pentland, A. (eds.) Artificial Intelligence for Human Computing. LNCS, vol. 4451, pp. 339–358. Springer, Heidelberg (2007).
22. Si, M., Marsella, S.C., Pynadath, D.V.: THESPIAN: An Architecture for Interactive Pedagogical Drama. In: Proc. 2005 conference on Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology. pp. 595–602. IOS Press, Amsterdam, The Netherlands, The Netherlands (2005).
23. Millington, I.: Artificial Intelligence for Games. Morgan Kaufmann, San Francisco (2006).
24. Niehaus, J., Weyhrauch, P.: An Architecture for Collaborative Human/AI Control of Interactive Characters. In: Proc. Workshop for Collaborative Human/AI Control for Interactive Experiences at Autonomous Agents and MultiAgent Systems (AAMAS). pp. 12-16. ACM, New York (2010).
25. Zhu, J., Moshell, J.M., Ontañón, S., Erbiceanu, E., Hughes, C.E.: Why Can't a Virtual Character Be More Like a Human: A Mixed-Initiative Approach to Believable Agents. In: Shumaker, R. (ed.) Virtual and Mixed Reality, Part II, HCII 2011, LNCS, vol. 6774, pp. 289-296. Springer, Heidelberg (2011).