# Creative Arcs in Improvised Human-Computer Embodied Performances

Mikhail Jacob
Georgia Institute of Technology
Atlanta, Georgia
mikhail.jacob@gatech.edu

Brian Magerko
Georgia Institute of Technology
Atlanta, Georgia
magerko@gatech.edu

## ABSTRACT

This paper proposes a novel framework for real-time decision-making in open-ended improvised human-computer performances and a work in progress system for studying it within the *Props* game domain. The proposed work uses creative arcs (i.e. continuous trajectories through a dimensional space consisting of novelty, surprise and value) as a way to guide the exploration of generated candidate actions for the improvising agent to perform. This paper also describes the proposed agent's relationship to curious agents and several considerations for computationally evaluating creativity in this domain. It details the interactive installation housing the system and the architecture design (completed and planned) before concluding with a discussion of future work.

## CCS CONCEPTS

• **Computing methodologies → Heuristic function construction**; *Cognitive science*; *Learning latent representations*; • **Applied computing → Performing arts**;

## KEYWORDS

Creative arc, open-ended improvisation, computational creativity, novelty, surprise, value

**ACM Reference Format:**
Mikhail Jacob and Brian Magerko. 2018. Creative Arcs in Improvised Human-Computer Embodied Performances. In *Foundations of Digital Games 2018 (FDG18), August 7–10, 2018, Malmö, Sweden.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3235765.3235827

## 1 INTRODUCTION

Improvisation with people is an important skill for creative agents to develop in order to act in realistically open-ended scenarios with massive knowledge/content requirements, cognitive/physical resource limitations, severe time constraints, a large number of potentially applicable actions, and the lack of a single clear goal to pursue at any time being characteristic of the problem. Various attempts have been made to solve constrained versions of this problem across improvisational domains including musical improvisation [12], improv theater [24], collaborative sketching [6], and improvised dance [14]. In contrast to previous approaches in this space using rule-based or stochastic processes to control real-time response generation, this paper proposes a novel framework for performing real-time decision-making in open-ended improvised human-computer performances using creative arcs (i.e. continuous trajectories through a dimensional space consisting of novelty, surprise and value) as a way to guide the exploration of generated candidate actions for the improvising agent to perform.

The goals for an agent to pursue at any given time in an open-ended improvisational performance (*improv* for short) are often poorly defined. It is also hard to provide objective fitness functions for evolving appropriate solutions due to the complexity of the improvisation task and the real-time nature of the improvisation. Additionally, reinforcement learning (RL) and imitation learning (from an RL perspective) are hard to use due to the lack of a suitable reward function and the massive, open-ended action space for the task respectively. We propose that this problem can be addressed by taking inspiration from the aesthetic trajectories used to both analyze and structure creative artifacts (and performances) in many distinct creative domains ranging from narrative to music to visual art. For example, agents in interactive narrative domains in the past [8, 21] have utilized Aristotelian dramatic arcs to guide action selection and improve the user's experience of the developing narrative. Music composition agents have used models of musical tension and release to structure their compositions [7]. This is similar to (but the inverse of) how narratologists and authors like Vonnegut have analyzed the shape or arc of stories [27]. Further, it is proposed that this approach can be generalized to guide action selection in open-ended improvisational domains.

This work proposes specifically that improvisational agents operating in open-ended creative domains over some time can select actions following a given 'creative arc' in a creative space in order to structure the improvised performance and improve a human collaborator's (or the audience's) subjective experience of the performance. This creative arc would either be given to the agent or be arrived at through some intrinsic motivational process (see **section 2**). The term creativity can be operationalized using Boden's definition of creativity [2] as the perceived novelty, surprise, and value of a creative artifact (see **section 3)**. An agent can thus choose actions to fit a given curve in the three-dimensional space of novelty, surprise, and value (i.e. a creative arc). Note that the creative space is too vast to systematically explore in a reliably reasonable time during improvised performance using random or brute force search alone. Therefore, improvisational agents can use response strategies adapted from human improvisers in other domains to guide this exploration of the creative space. Thus the proposed agent has at least four components:


*FDG18, August 7–10, 2018, Malmö, Sweden*

ACM ISBN 978-1-4503-6571-0/18/08…$15.00
https://doi.org/10.1145/3235765.3235827

(1) a *parameterized generator* for systematically exploring the space of candidate actions that an agent can take
(2) a set of computational models for *evaluating* the *novelty*, *surprise*, and *value* of candidate actions
(3) a set of *response strategies* adapted from human improvisers to bias the agent's exploration towards obtaining desirable candidates as quickly as possible
(4) an *adaptive controller* to select from the response strategies in order to follow the given creative arc.

The long-term goal of this research is to use the proposed decision-making framework to enable embodied virtual characters to participate in embodied narrative improvisation with people in real-time [13]. Embodied narrative improvisation is applicable to game AI, human-agent interaction, human-robot interaction, and other situations where the agent performs creative behaviors using its bodily capabilities to interact with the environment in which it is situated. Embodied narrative improvisation with people in an open-ended problem domain involves narrative intelligence, social cognition, linguistic and non-linguistic action, and other cognitive faculties.

As a first step in the direction of embodied narrative improvisation, a constrained domain was chosen for study. The 'Props' game is a short form improv theater game that involves improvised interactions between two or more participants using unfamiliar ambiguous props to perform recognizable comedic actions pretending the prop to be a familiar real-world object. In this case, the improv will take place between an embodied virtual agent and its human collaborators using ambiguous props that are potentially unfamiliar to the agent. Beyond improv theater, this problem is useful for embodied agents in general since it is the first step towards allowing them to gain new knowledge about unfamiliar objects through interaction. For example, this could include an agent learning to use unfamiliar objects in unfamiliar scenarios according to familiar human norms/customs or using unfamiliar objects for a specific task (such as improvising a digging tool for disaster recovery).

The remainder of this paper describes how creative arcs relate to curiosity (as well as other models of intrinsic motivation) for intelligent agents and useful models (and special considerations) for computationally evaluating creativity in this domain. It then describes a forthcoming public virtual reality (VR) installation that will be a test bed and technical probe for studying the creative arc following agent architecture. The installation (called the *Robot Improv Circus*) will enable participants to play the Props game with a virtual character (or watch others do so) in real-time. The paper then describes the design of the work-in-progress agent architecture (called CARNIVAL) proposed earlier. Finally, the article concludes by discussing future work remaining.

## 2 CREATIVE ARCS, CURIOSITY, AND INTRINSIC MOTIVATION

Creative arcs (as defined above) are continuous trajectories through a dimensional creative space that an agent can aim to follow over the course of a temporally extended improvised performance. For example, one such creative arc could have the agent perform relatively low novelty, low surprise, and high value actions at the start of the performance, rising up to high novelty, high surprise, and high value actions towards the middle two thirds of the performance, before finally culminating in medium novelty, medium surprise, high value actions towards the end of the performance. Some examples of creative arcs are illustrated in **Figure 1**.

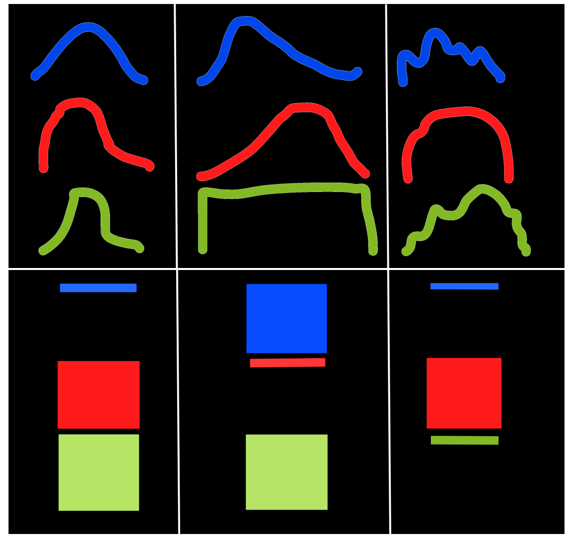

**Figure 1: A). A series of example curves in the novelty (blue, top), surprise (red, middle), and value (green, bottom) dimensions. B). Three example curves that simulate a novelty searching agent (left) with maximized novelty, a surprise searching agent (middle) with maximized surprise, and a hybrid agent [22] (right) with maximized novelty and value. The filled box represents that the entire region is acceptable (effectively ignoring it).**

The system is not restricted to following any particular creative arc. They can be given to the agent by the creator of an improvised performance (or any external source). This is one way that the agent can personalize the experience to a particular collaborator or audience or situation. Alternatively, the agent itself could generate different creative arcs, searching over them at a meta-level across many performances to learn which ones work better than others.

Curiosity — as used to describe an intrinsic motivating drive to discover novel percepts, experiences, or knowledge — has been used in techniques such as novelty search [18] to discover better solutions compared to pure objective search [18]. We state that it is possible to simulate a similar novelty seeking (curious) agent in the proposed decision-making model by providing it with a creative arc that has a uniformly maximal novelty dimension along with setting the agent to ignore (or alternatively, to accept any value in) the surprise and value dimensions of the creative space (see **Figure 1**). This also enables the agent to perform surprise search [10] by providing the system with a creative arc that maximizes the surprise dimension while ignoring (or alternatively accepting any value in) the other two dimensions. Additionally, it can also simulate other hybrid search agents [22].

The creative arc following agent proposed in this work differs from the various intrinsically motivated agents mentioned above in the following two ways. Firstly, the former directly optimize novelty, surprise, and value dimensions (among others) while the latter tries to optimize a given meta-level function composed of those dimensions. Therefore, a creative arc that starts with low novelty and progresses to some peak novelty value before descending again might be more valuable to an improvisational partner than an agent that tries to do the most novel action it possibly can every single turn. Secondly, in the former case there is often a final output to the search process (when search is stopped eventually) that is evaluated to assess the effectiveness and quality of the optimization technique, while in the latter case, the agent's creative artifacts are experienced by the agent's improvisational partner (and a potential audience) all throughout the creative arc making the journey itself the main creative artifact that is assessed and not necessarily any individual action generated by the agent along the way.

## 3 COMPUTATIONAL EVALUATION OF CREATIVITY

There have historically been an extraordinary number of models of creativity and the creative process from diverse research fields ranging from media studies to psychology to artificial intelligence and computational creativity. However, the following have been most useful for deriving a computational model of creativity evaluation in our work. Newell et al.'s model [23] for evaluating creative problem solving referred to novelty, value, rejection of previous assumptions, persistence towards a goal, and development of the problem specification itself. Boden's influential model [2] of creativity involved evaluating the novelty, surprise, and value of the generated artifacts. Maher [20] operationalized Boden's model of creativity for evaluating creative artifacts. Colton's creativity tripod [4] argued for a computationally creative system having skill, imagination, and an appreciation of the creative medium. Colton et al.'s FACE model [5] evaluated computational creativity systems for creative concept invention, expression of the concept as an artifact, aesthetic evaluation of the artifact, and the framing of the artifact to the public. Finally, Jordanous' Standardized Procedure for Evaluating Creative Systems (SPECS) methodology [15] advocated for creating a customized working definition of creativity for the computational creativity system and then adapting experiments or evaluations tailored to that working definition. All of the above, except Maher [20], provided non-computational frameworks or methodology for evaluating creativity. Therefore, our work extends Maher [20] to provide computational models for evaluating creativity for improvisational domains as detailed below.

The working definition of creativity used in this work is an extension of Boden's definition of creativity focusing on the novelty, surprise, and value of generated artifacts. We propose a multidimensional model of creativity, with an artifact localized to a point in the space of novelty, surprise, and value. In addition, the proposed set of models used to localize candidate actions in the creative space potentially differ from other models in the literature along the following dimensions due to the specific constraints of open-ended improvisation between a human and virtual character (with a potential audience). These constraints are listed below.

(1) The *perspective* being evaluated: For an improvised performance/interaction between a virtual character, human collaborator, and an audience there are three separate perspectives for judging the creativity of the improvised interactions. The choice would depend on the main goal of the interaction, whether to optimize the quality of agent's learning & data acquisition, the user experience of the human collaborator, audience enjoyment, or some combination of these (ideal for an improvised human-computer performance).
(2) The degree of *dynamism*: This is the amount that the evaluation changes over time due to the experiences of the agent. A static/unchanging model would be fixed (without accounting for habituation or other changes over time), while a more dynamic model might adapt offline in between every improvisational session. The most desirable model adapts online over the course of the ongoing improvisational session.
(3) The role of *feedback*: The model may not use feedback at all to improve its scoring over time. Alternatively, the model might utilize explicit feedback from the audience (e.g. applause) or collaborator (e.g. post-interaction surveys). The feedback could also be implicit through metrics like interaction duration or facial expression counts if explicit feedback can't easily be collected. Feedback is usually desirable unless the expertise of the system is far greater than the user.
(4) The relative *expertise* of the system: A fledgling system that has little data or experience cannot expect to match human ratings of novelty and expectation and should treat the user's experiences as a superset of its own (e.g. an open-ended narrative improv system). A system that has collected data over its lifetime or through massive datasets can potentially surpass the human in terms of experience (e.g. a recipe generation system mining from large online recipe databases). It might then need to localize novelty and surprise estimation to the neighborhood of the user's experiences.

## 4 THE ROBOT IMPROV CIRCUS

The Robot Improv Circus is a VR installation for a single participant to play the Props game on stage in a virtual circus where everyone is a humanoid robot (the participant, their virtual stage partner, the virtual audience, and the virtual judges). **Figure 2** shows prototype versions of the physical installation and the virtual world. While the participant is interacting with the virtual world, the real-world audience can also participate by providing feedback to the participant and the virtual improviser through applause, by viewing their performance in the virtual world, or by directly rating the ongoing performance. The installation serves as a test bed and technical probe for studying improv between humans and virtual agents using human-centered methodology.

## 5 THE CARNIVAL ARCHITECTURE

The CARNIVAL (Creative ARc Negotiating Improvisational Virtual Agent pLatform) intelligent agent architecture for performing open-ended embodied improvisation with people consists of four components. A deep generative model generates candidate actions for the agent to perform. A set of heuristic models evaluate candidate actions to predict their coordinates in a formal creative space. A set of improvisational response strategies adapted from human improvisers is applied to the generative model to guide its search


**Figure 2: A). An initial outer design for the Robot Improv Circus installation. B). An initial prototype of the Robot Improv Circus VR experience where humans and robots can play the Props game together.**

through the generated action space in order to find suitable candidate actions for the agent to perform in real-time. Finally, given a desired creative arc for the performance, an adaptive controller selects strategies in order to generate actions based on the agent's target position on the creative arc throughout the performance.

### 5.1 Deep IMprovised Action Generation through INteractive Affordance-based exploraTION (DeepIMAGINATION)

The approach used in this system to improvise open-ended performances involves guided search over a possibility space, evaluating candidate responses that the agent could perform. This requires a systematic parameterized generator for the target domain. The candidate responses produced by the generator are then evaluated by other evaluation components for suitability or fitness.

The specific version of this problem in the problem domain is the parameterized generation of valid mimed actions for the agent using a given ambiguous prop. The DeepIMAGINATION (Deep IMprovised Action Generation through INteractive Affordance-based exploraTION) module in the agent is responsible for generating candidate actions accordingly. DeepIMAGINATION uses two insights to solve this problem. Firstly, since the props are ambiguous enough to be imagined as multiple real-world objects, it represents and reasons about the props more generally using their physical attributes or affordances. Secondly, it formalizes the theoretical process of searching through an action space as sampling from the latent space of a conditional variational autoencoder (CVAE) [26] conditioned on the physical affordances of props.

The physical affordances of the prop are represented as a feature vector representing aggregated counts of features for all part of the prop summed together (e.g. how many spherical, cylindrical, thin, thick, or tapered parts are present in the prop). Each part is a component of the prop that corresponds to a shape primitive with (optional) deformations applied to it. For example, a barbell-shaped prop might be two flattened spheres and a long, thin cylinder connecting them. This encoding is done by hand (for now) given the limited number of props currently in use.

The initial version of the conditional variational autoencoder (CVAE) [26] is trained on novice improvised mimed gestural actions using the props as pretend real-world objects within a VR data collection environment. Training input consists of a vector of location data for the user's head and hands (and the character's pelvis) along with the two controllers' prop grab/drop button states, sampled at 90 FPS for 10 seconds (with zero padding for shorter actions). The encoder and decoder are both conditioned on the encoding of the physical affordances of the props used to perform the input actions. They use convolutional layers (and transposed convolutional layers respectively) with dropout to encode the high dimensional input (27000-dimensional vector) into a low dimensional latent space (multiple versions have been tested using different numbers of latent variables from 2 to 16) and to decode it back again.

The network is trained by minimizing the variational autoencoder loss function (sum of decoder's reconstruction loss and encoder's Kullback-Leibler divergence regularization loss) using the re-parameterization trick from Kingma and Welling [16]. During generation, the latent space is sampled as desired to generate actions conditioned on the physical affordances of the given prop. The generated output is then used within the Robot Improv Circus VR experience to control the rigged character model of a virtual character using inverse kinematics (to control the character's other skeletal joints). Smoothing of generated joint trajectories is necessary due to noise in the output resulting in shaky movements.

### 5.2 Creative Space Evaluation Heuristics

While the DeepIMAGINATION module can generate (potentially) valid actions for a given prop, the agent's creative process requires a set of criteria for selecting the agent's response during its turn. This is particularly challenging when doing open-ended improv as there are not specific, predefined goals for the agent to satisfy, but rather high-level heuristics that suggest that some responses are better than others. Additionally, since the decision-making of the agent is motivated not by concrete goals, but by following a creative arc (see **section 2**), the agent requires some way to localize where actions are located in the creative space with respect to the desired creative arc. This is accomplished in the present system using a set of evaluation heuristics to predict the novelty, surprise, and value of the agent's and collaborator's actions.

The creative space evaluation heuristics are a work in progress with initial prototyping under way. They consist of heuristics for evaluating the predicted novelty, surprise, and value of the generated actions extending previous work in the literature. These heuristics are described briefly below.

*5.2.1 Novelty.* Novelty is defined here as the degree to which some percept/experience is different to an agent compared to other comparable percepts/experiences the agent has seen before, bounded by some experiential time horizon. In the current system, novelty is measured as the aggregated distance between a generated/observed artifact and the nearest **K** action clusters in the latent space. As an initial simplification, the novelty heuristic currently only considers the agent's perspective and does not consider feedback from others to improve over time. The current model being implemented does include dynamics over the course of the ongoing interaction session (e.g. habituation as the agent experiences multiple similarly novel percepts/experiences). The current model is expected to be a novice compared to human experience in the domain, but that would change as the system learns over its lifetime.

*5.2.2 Surprise.* Surprise is an affective reaction to the violation of a confidently held expectation and is proportional to the degree of expectation violation or unexpectedness [9]. Surprise presents an opportunity to investigate why that expectation was violated and to refine the model or reasoning that produced the violated expectation [25]. Surprise has been measured in various ways in the past depending on the specific context. The approach currently being prototyped in this system considers the DeepIMAGINATION module to hold confident expectations about action generation with the given prop, therefore the degree of deviation from expected actions is measured as the agent's surprise. For this initial prototype, the model does not consider other perspectives and does not consider feedback. However, it displays dynamism as the system learns more over time (though DeepIMAGINATION is only retrained in between sessions) and the system's expertise is closer to a novice though that would change over time.

*5.2.3 Value.* Value is the most open-ended and least well-defined of the measures described in this work. Value is also one of the components of creativity that is extremely domain dependent and can be decomposed into several component heuristics according to the domain. The total value score would be composed of a combination of these individual component scores individually weighted according to empirical evaluation. A set of proposed components for the value function are outlined below.

One important class of component heuristics measuring value in this domain could be the kinds and intensity of affect produced by a candidate action. For example, measuring the humor of a candidate action or the pathos of a candidate action could be two different components of value. However, it might be that humor is generally more important since the domain is improv theater. Computational models of humor relying on measures of benevolent expectation violation could be particularly applicable in this system since unexpectedness and surprise are already measured.

Another important heuristic is the recognizability of the pretend object or to what degree the mimed action with the ambiguous prop correctly signifies the intended pretend object. The complementary value heuristic is the recognizability of the mimed action or to what degree the mimed action with the ambiguous prop can be correctly interpreted as the intended action. These are related to iconicity [19] which indicates how relatively unique the value of an attribute is to a category in order to signify membership to it.

The value of the improvised performance could also be evaluated in terms of the aesthetic value of the performed actions. For example, measuring the smoothness of the paths/trajectories taken by the virtual character's limbs could measure the aesthetics of their limb motion. Other heuristics from movement theory could also be used to develop measures of aesthetic value such as how expressive a motion is in one of the Viewpoints [3] or Laban [17] dimensions or how well the motion adhered to Bartenieff Fundamentals [1].

Another measure of value in the context of the improvisation could be the degree of connection to the collaborator's last mimed action or pretend object. For example, pretending the current object is a shield when the collaborator pretended on their last turn that they were swinging a sword could have higher value than a completely disjointed mimed action and pretend object on the agent's turn. Related measures of value could be the use of reincorporation to refer to collaborator mimed actions or pretend objects from the past or the use of callbacks for repeating a particular mimed action or pretend object later in a different (but appropriate) context.

## 5.3 Response Strategies

Response strategies for guiding the search through the creative space are intended to enable the agent to quickly arrive at acceptable points in the creative space. These response strategies were adapted from the kinds of high-level strategies commonly used by jazz improvisers [11]. They were previously used in the LuminAI system as well [14]. An incomplete list of some common strategies follows.

Mimicry (i.e. repetition) occurs when the agent chooses the same mimed action with the same pretend object that their collaborator did. This will be implemented in the current system by obtaining the latent space encoding of the collaborator's last action and generating the action from that latent space coordinate through the decoder. In previous systems [14], this response strategy created outputs that were generally lower in novelty, but relatively higher in value depending on how often it was used.

Transformations (i.e. augmentations) of various types take the collaborator's last mimed action and apply modifications to it in order to inject some novelty into the action while maintaining a strong connection to the collaborator's last creative offer. Some examples include, choosing a similar but related (and valid) pretend object for the same action, using the generator to output actions from similar but non-identical latent space coordinates, and affine transformations of the line connecting the last two actions in the latent space (this would be similar to an analogical transfer of the latent space relationship between the last two actions of the improvisation to a new part of the latent space). In previous systems, transformations created outputs that exhibited varying degrees of novelty or surprise depending on the transformation used.

Novel action generation, takes the form of random generation of an action from some region of the latent space or generation from a different region of the latent space compared to where the agent currently is exploring (e.g. intentionally jumping to a new region of the latent space). Random generation is a wild card for affecting the evaluation of generated artifacts predictably. However, the latter strategy could be a novelty-increasing operation.

Combination refers to the interpolation between various numbers of candidate actions by taking two or more points in the latent space and using an intermediate point such as a weighted centroid to interpolate between them in the latent space. This strategy has the potential for increased surprise and novelty. However, there is a higher chance of generating artifacts of lower value.

Value-modulation strategies are those operations that are tailored towards increasing or decreasing the value of a given artifact. For example, intentionally generating a regular action modified so that the prop is dropped halfway through the action for humor is an extradiegetic increase of the humor value of the action. Alternatively, performing a callback by repeating the collaborator's high value mimed action from a few turns ago when the same prop is given to the agent later is a diegetic value increase.

### 5.4 Strategy Selection Controller

The final component required for the system is an adaptive controller for selecting which response strategy the agent should use at any point. The controller receives the desired creative arc and the desired session duration at the start of the session. Then for each of the agent's turns, it chooses the response strategy for that turn by considering the vector connecting the location of the agent's last generated action in the creative space and the target location in the creative space for the agent's current turn according to the given creative arc. According to the desired direction in each dimension of the creative space, the controller selects the response strategy that it predicts will take it nearest to the target location. For each strategy, the controller has a certain confidence threshold that changes based on the success of recent usage results. This is also considered when choosing the response strategy.

## 6 CONCLUSION

This paper contributes a novel framework for real-time decision-making in open-ended improvised human-computer performances and a work in progress system for studying it within the Props game domain. The proposed work uses creative arcs (or continuous trajectories) through a dimensional space consisting of novelty, surprise, and value as a way to guide the exploration of generated candidate actions for the improvising agent to perform. The paper also describes the creative arc following agent's relationship to curious agents and some other intrinsically motivated agents. Several considerations for computationally evaluating creativity in this specific domain are also outlined. This paper describes both the Robot Improv Circus interactive installation and the CARNIVAL intelligent agent architecture controlling its virtual improviser.

From the proposed CARNIVAL architecture, the DeepIMAGINATION module has been completed and is now being formally evaluated. The creativity evaluation heuristics are currently being developed in parallel. The response strategies and the adaptive controller will be developed by the end of the year. Throughout the year, each component will be iteratively evaluated and refined.

Future study will focus on understanding the effect of the decision-making approach on various user experience metrics compared to other baseline techniques. Additionally, while the current CVAE generator used in DeepIMAGINATION is restricted to locally exploring candidate actions in the neighborhood of those that have been seen in human data, it would be invaluable to be able to generate completely novel actions that might still be valid for a given prop. The creativity heuristics will also be improved by considering multiple perspectives and interaction feedback into account. Finally, the agent currently attempts to follow a given creative arc. However, in the future it could also internally generate creative arcs for different scenarios and different individuals.

## REFERENCES

[1] Irmgard Bartenieff and Dori Lewis. 2013. *Body movement: Coping with the environment.* Routledge.
[2] Margaret A Boden. 2003. *The creative mind: myths and mechanisms.* Psychology Press. DOI: 10.1017/S0140525X0003569X.
[3] Anne Bogart and Tina Landau. 2006. *The Viewpoints Book: A Practical Guide to Viewpoints and Composition.* Theatre Communications Group.
[4] Simon Colton. 2008. Creativity Versus the Perception of Creativity in Computational Systems. In *Proceedings of the AAAI Spring Symposium on Creative Systems.* 14–20.
[5] Simon Colton, John William Charnley, and Alison Pease. 2011. Computational Creativity Theory: The FACE and IDEA Descriptive Models.. In *Proceedings of the 2nd International Conference on Computational Creativity.* 90–95.
[6] Nicholas Davis, Chih-PIn Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces.* ACM, 196–207.
[7] Ramon Lopez De Mantaras and Josep Lluis Arcos. 2002. AI and music: From composition to expressive performance. *AI Magazine* 23, 3 (2002), 43.
[8] Magy Seif El-Nasr. 2007. Interaction, narrative, and drama: Creating an adaptive interactive narrative using performance arts theories. *Interaction Studies* 8, 2 (2007), 209–240.
[9] Kazjon Grace, Mary Lou Maher, Douglas Fisher, and Katherine Brady. 2015. Data-intensive evaluation of design creativity using novelty, value, and surprise. *International Journal of Design Creativity and Innovation* 3, 3-4 (2015), 125–147.
[10] Daniele Gravina, Antonios Liapis, and Georgios Yannakakis. 2016. Surprise search: Beyond objectives and novelty. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference.* ACM, 677–684.
[11] Robert Hodson. 2007. *Interaction, improvisation, and interplay in jazz.* Routledge.
[12] Guy Hoffman and G Weinberg. 2010. Shimon: an interactive improvisational robotic marimba player. *CHI'10 Extended Abstracts on Human Factors \ldots* (2010), 3097–3102.
[13] Mikhail Jacob. 2017. Towards Lifelong Interactive Learning For Open-ended Embodied Narrative Improvisation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition.* ACM, 502–507.
[14] Mikhail Jacob and Brian Magerko. 2015. Interaction-based Authoring for Scalable Co-creative Agents. In *Proceedings of the Sixth International Conference on Computational Creativity (ICCC 2015).* Provo, UT.
[15] Anna Katerina Jordanous. 2011. Evaluating Computational Creativity: A Standardised Procedure for Evaluating Creative Systems and its Application. (2011).
[16] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
[17] R. Laban and L. Ullmann. 1971. The mastery of movement. (1971).
[18] Joel Lehman and Kenneth O Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011), 189–223.
[19] Brian Magerko, Peter Dohogne, and Chris DeLeon. 2011. Employing Fuzzy Concepts for Digital Improvisational Theatre. *AIIDE* (2011).
[20] Mary Lou Maher. 2010. Evaluating creativity in humans, computers, and collectively intelligent systems. In *Proceedings of the 1st DESIRE Network Conference on Creativity and Innovation in Design.* Desire Network, 22–28.
[21] Michael Mateas and Andrew Stern. 2003. Facade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference, Game Design track*, Vol. 2. 82.
[22] Jean-Baptiste Mouret. 2011. Novelty-based multiobjectivization. In *New horizons in evolutionary robotics.* Springer, 139–154.
[23] Allen Newell, J Clifford Shaw, and Herbert Alexander Simon. 1959. *The processes of creative thinking.* Rand Corporation Santa Monica, CA.
[24] Brian O'Neill, Andreya Piplica, Daniel Fuller, and Brian Magerko. 2011. A knowledge-based framework for the collaborative improvisation of scene introductions. In *Proceedings of the 4th International Conference on Interactive Digital Storytelling*, Vol. 7069 LNCS. Vancouver, Canada, 85–96.
[25] Wolfram Schultz, Peter Dayan, and P Read Montague. 1997. A neural substrate of prediction and reward. *Science* 275, 5306 (1997), 1593–1599.
[26] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems.* 3483–3491.
[27] Kurt Vonnegut. 2005. At the Blackboard. *Lapham's Quaterly* (2005).