Empirically Studying Participatory Sense-Making in Abstract Drawing with a Co-Creative Cognitive Agent

Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, Brian Magerko
Georgia Institute of Technology
School of Interactive Computing
{ndaivs35, chsiao9, kysingh, lisa_li, magerko}@gatech.edu

ABSTRACT

This paper reports on the design and evaluation of a co-creative drawing partner called the Drawing Apprentice, which was designed to improvise and collaborate on abstract sketches with users in real time. The system qualifies as a new genre of creative technologies termed "casual creators" that are meant to creatively engage users and provide enjoyable creative experiences rather than necessarily helping users make a higher quality creative product. We introduce the conceptual framework of participatory sense-making and describe how it can help model and understand open-ended collaboration. We report the results of a user study comparing human-human collaboration to human-computer collaboration using the Drawing Apprentice system. Based on insights from the user study, we present a set of design recommendations for co-creative agents.

Author Keywords

Creativity Support Tools; Computational Creativity; Collaboration

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

INTRODUCTION

Collaboration is a powerful way to help both novices and experts push their creative boundaries and inspire their creative process. As collaborators make complimentary and unexpected contributions, their shared product grows in an emergent manner that would be more creative than what each partner could achieve individually [30]. Developing agents to facilitate collaboration can provide insight into the nature of collaboration as well as means of supporting collaborative creativity. To investigate improvisational collaboration, we developed a co-creative drawing partner called the Drawing Apprentice, shown in Figure 1. This
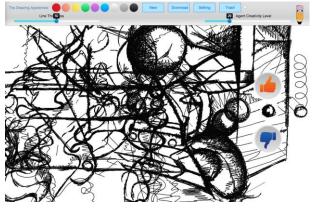
Figure 1: Drawing Apprentice UI and collaborative drawing made with the system.

system collaborates with users in real time using turn taking to gradually build a drawing through time [10].

We employ the cognitive science theory of enaction [34] and its conceptual framework called *participatory sense-making* [15] to model and understand creative collaboration. This framework describes critically important cognitive mechanisms and processes that we have seen evidence of in other collaboration tasks, such as joint activity and cooperation [7,13]. Participatory sense-making has been utilized to characterize collaboration and narrative development in pretend play [7], which is a similarly improvisational and open-ended domain. We used participatory sense-making to inform our system design and analyze improvisational collaboration in the domain of abstract drawing.

Cognitive agents, from the perspective of enaction, interact to gradually determine patterns of regularity and meaning using dynamic feedback loops [15]. Therefore, each collaborator is influenced both by the content of actions (i.e. lines of a drawing) as well as the dynamics of interaction that have emerged during in-the-moment interaction, such as the rhythm of turn taking, area of focus, and manner of motion. This dual feedback system has the potential to lead collaborators into new realms of creative expression with dynamic, multi-layered, and sometimes competing mental models for evaluating what creative contributions 'make sense' in the current situation.

Given the novelty and complexity of open-ended creativity and collaboration as a research topic, our prototype serves

primarily as a technical probe to identify the circumstances under which contributions from co-creative agents appear to make-sense. In particular, we investigate three dimensions of participatory sense-making in collaborative drawing:

- To what degree was participatory sense-making present during the collaboration?
- What metrics and features did users employ to determine whether contributions 'made sense'?
- How did users try to define shared meaning structures with the agent, i.e. how did they attempt to teach the system?

We report on a Wizard of Oz-style user study comparing human-human collaboration to human-computer collaboration with the Drawing Apprentice system. We hypothesize that a system that mimics and builds on the contributions of users using real-time feedback can impact the sense-making process of art making with similar benefits as participatory sense-making exhibited in open-ended creative collaboration between two people. Our findings show that the Drawing Apprentice was able to engage users in participatory sense-making that resulted in discovering novel visual ideas and emergent meaning in the drawing. We also present several metrics users employ when evaluating whether contributions from the agent 'made sense.' These findings are then summarized in a set of design recommendations for supporting participatory sense-making with co-creative agents more generally.

RELATED WORK

The work reported here is an interdisciplinary effort nestled between the fields of computational creativity and creativity support tools. Computational creativity is a field of artificial intelligence focused on developing agents that generate creative products autonomously [3,5,38,39]. Creativity support tools, on the other hand, are technologies designed to enhance and augment the user's creativity, typically aiming to improve the quality of the final product [12,32,33,35]. Computers can support human users in their creative process in a variety of ways, including acting as a nanny, coach, pen-pal, and colleague [26]. The Drawing Apprentice can be considered a computer colleague since it collaborates with users on their creative task.

There are some computational creativity projects in other creative domains that explore human-computer co-creativity, such as Weinberg's work in interactive improvisational percussion robots [21]. Weinberg's robot, Shimon, listens to a performer and mimics or adds to the performance of the human collaborator. The system analyzes the music of performers and generates melodic improvisations that are in sync with human collaborators. In practice, the human and robot develop a call and response interaction where each party modifies and builds on the previous contribution. GenJam is another improvisational jazz system that uses a form of genetic algorithms to generate jazz improvisations [2]. The system has acoustic sensors that recognize musical input. A number of jazz improvisation schemes are stored in the program. An accompaniment is selected by the computer and played on a MIDI style system. Biles has performed with GenJam in jazz clubs successfully for several years. He advocates evaluating such systems based on the degree to which the system is useful and meaningful to the artist.

Viewpoints AI (VAI) is a co-creative dance partner that improvises with users in real time as they dance in front of a virtual character projected on a large display screen [22]. Using a Kinect, VAI analyzes the dance gestures of users and selects a complimentary dance move for the virtual character to perform. The system was initially trained by analyzing professional dancers to seed its knowledge base with expert dance moves [23]. Its knowledge base grows as it observes novel moves performed by users. VAI has a similar goal as the Drawing Apprentice to coordinate with users during open-ended creative improvisation. However, drawing results in a creative product that remains visually present and grows over time versus the ephemeral activity of dancing.

Another example of human-computer co-creativity is Magerko et al.'s Digital Improv project [27–29]. In this project, users engage in multiple theatrical improvisation games focused on narrative construction with AI improvisers. In the Digital Improv project, a computationally creative system attempts to recognize actors' actions, interpret them to be relevant to kinds of characters, and respond to these actions in a virtual world. The computer develops a shared meaning with the actors and generates actions according to features of the negotiated narrative context [17–20].

In the domain of artistic creativity, there are a variety of CSTs focused on supporting artistic creativity in the capacity of a coach to train the user's technical skill. For example, ShadowDraw [25] attempts to recognize the object users are drawing and provide them with guidelines in real-time to help draw more accurate representations. iCanDraw [11] helps users draw portraits by providing feedback on how well the user's lines align with the target portrait to enhance the accuracy of drawing their portraits. Projector Guided Painting helps users reproduce classic artworks by using a projector to depict the next color that should be applied to a particular region of the canvas [14].

These CSTs focus on increasing artistic skill and helping users create high quality products, but individuals interested in learning to engage in creative expression for its own intrinsic rewards have unique needs that often fall outside of existing CSTs. These users are not necessarily focused on achieving the best product, but derive enjoyment from the creative process itself, i.e. *autotelic creativity*, rather than task-oriented creativity [6]. A new category of tools was recently recognized that are primarily designed to creatively engage these casual users called *casual creators*, defined as:

"...an interactive system that encourages the fast, confident, and pleasurable exploration of a possibility space, resulting in the creation or discovery of surprising new artifacts that bring feelings of pride, ownership, and creativity to the users that make them." [5]

Casual creators are systems designed specifically to support the process and experience of creativity. In this context, the fun and enjoyable dimensions of the actual creative experience interacting with the tool are emphasized over the aesthetic quality of the final product.

Potential use cases for the Drawing Apprentice include engaging children with the tool to facilitate creative 'out of the box' thinking by encouraging them to integrate and adapt to novel artistic contributions, thereby increasing their visual lexicon and ability to integrate unexpected contributions. In particular, children on the autistic spectrum may derive unique benefit from the tool by encouraging joint activity and adaptation, which are cognitive skills with which this population traditionally struggles [4,31].

PARTICIPATORY SENSE-MAKING

There are five core ideas to the cognitive science theory of enaction that are helpful in considering the design of a co-creative improvisational agent. First is the notion of *autonomy*, which describes how cognitive agents operate independently, based on their own intrinsic laws, to sustain themselves and continually generate their identity through interaction with the environment according to those laws. A system is defined as 'autonomous' when it can change the laws that govern its interaction with the world by casting a web of significance on the elements in its environment. As autonomous agents interact with their environment, they gradually detect patterns of regularities from dynamic feedback loops, which help them make sense of the environment in a process referred to as *sense-making*. Enaction describes the *emergence* of meaning structures through coordination and coupling between a cognitive agent and some system in the environment. It also adheres to the *embodiment* paradigm in asserting that agents must act with their body to make sense of the environment, which inherently constrains and affords certain types of interaction. Finally, the subjective *experience* of the agent is shaped by the meaning it has imbued into its environment and this influences how future actions are generated and evaluated.

A cognitive agent can be said to exhibit a degree of *structural coupling* with a system in the environment when it has understood the parameters and mechanisms of that system to the extent that it can predict how different actions will affect the system. When additional agents are introduced into the fold, the dynamics of sense-making become more complex because each agent is independently engaging in a process of sense-making with the environment as well as each other. This mutually

influencing process is termed *participatory sense-making*, and it is formally defined as:

> A co-regulated coupling between at least two autonomous agents, where: (i) the co-regulation and the coupling mutually affect each other, constituting an autonomous self-sustaining organization in the domain of relational dynamics and (ii) the autonomy of the agents involved is not destroyed (although its scope can be augmented or reduced). [24]

This quote outlines the layered process of meaning making in participatory sense-making where the method for interacting together exists as an additional dimension, here described as relational dynamics, i.e. the rhythm of turn taking, ways of providing feedback, and manner of actions (independent from the content of actions).

For true co-regulation to occur in participatory sense-making, agents must make sense of both these relational or interaction dynamics as well as the content of actions in a dual sense-making process that is unique to participatory sense-making. We expect to find clear examples of such co-regulation in human collaboration, and through our experimentation, we hope to define the types of parameters and mechanisms that would enable a co-creative agent to exhibit the degree of autonomy and sense-making that would give rise to a similar co-regulated participatory sense-making as seen in human collaboration. This type of interaction is defined as the ideal for creative collaboration due to its ability to facilitate emergent meaning in an open-ended interaction that goes beyond what the individual user could have accomplished alone.

Enactivist researchers define cooperation as follows: "cooperation means that each of the subjects is taking account of the other's interests and objectives in some relation to the extrapersonal context, and is acting to complement the other's response [13]." We argue this definition is applicable to improvisational creative collaboration, and the means by which it can be achieved is through a process of co-regulated participatory sense-making. During this dialogical process, the dynamics of the interaction itself become autonomous and have the potential to propel the interaction forward. Interaction dynamics, in the case of digital drawing collaboration, include elements such as the length of a turn, the rhythm of actions within that turn, and the flow or trajectory of the experience, i.e. staying in one region versus moving to another.

We aim to explore with this work how a co-creative agent may be able to facilitate some degree of participatory sense-making without high-level representations of aesthetics and art theory. The quality of sense-making between a human collaboration and computer collaboration can then be compared to identify critical elements that were present in human collaboration, but absent when collaborating with the agent. We can then apply those insights to update the prototype and conduct further experimentation to evaluate whether users were able to achieve cooperation and collaboration in a more meaningful manner.
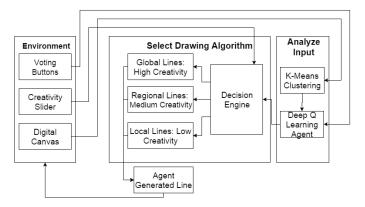
Figure 2: Drawing Apprentice System Architecture Diagram

DRAWING APPRENTICE SYSTEM

In this section, we describe the important technical components of the Drawing Apprentice system architecture and explain how they relate to the user experience. The system is implemented as a web application[1] with a client-server architecture that enables multiple people to collaborate with each other as well as the agent from the Drawing Apprentice system. It was designed for use with stylus- or touch-based interactions, but mouse can also be used.

At the highest level, the system is a co-creative agent that takes user input lines, transforms those lines based on pre-encoded line transformation techniques, and outputs new lines onto the same canvas (see Figure 2). Unique and defining features of input lines set are determined by clustering the data points in the input lines and sending that cluster data into the neural network (see Analyze Input box in Figure 2). This allows the neural network to derive its own classifications scheme based on the data it has been given.

The system was seeded with 12 experimental line transformational algorithms, including simple functions, such as translation, scaling, rotation, as well as more complex techniques that change the individual features of the input lines to create new lines that retain a similarity to the input lines. These more complex transformations are achieved by determining a set of equations to describe the input line and then tweaking individual coefficients in the equations to produce a similar line (see Figures 4 & 5).

Figure 3: The drawing results from the algorithm 1- 3 for low-creativity level (black lines: human, blue lines: agent)

Inspired by the theory of enaction, the system is designed to engage in participatory sense-making by coordinating its actions through real-time feedback. In practice, this means

[1] https://adam.cc.gatech.edu/DrawingApprentice/

training the system's machine learning algorithms dynamically using real-time feedback from the voting buttons and creativity slider.

The creativity slider in the interface of the Drawing Apprentice constrains which algorithms the system will choose from when reacting to the user's lines. The creativity level roughly corresponds to the general definition of the term 'creativity' in the creativity literature as: novelty, value, and surprise. At low levels of creativity (slider is between 0-33), the system will produce lines that slightly alter the user's input lines without much change. For example, in Figure 3-Algo-2, the system introduces some noise or perturbation to the input line and then redraws it. This type of contribution is almost identical to the user's input line and therefore is not very 'novel' or 'surprising.' Other algorithms at the low level of creativity mimic the user's line, but slightly translate or offset it, e.g. Algo-3 in Figure 3.
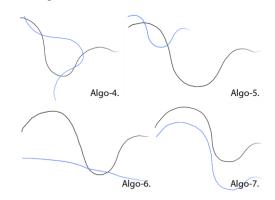

Figure 4: The drawing results from the algorithm 4-7 for mid-creativity level (black lines: human, blue lines: agent)

When the system is set to medium creativity (slider is situated between 33-66), the agent mimics and imitates the user's input with transformations, such as rotation (Algo-4 in Figure 4) and scaling (Algo-5 in Figure 4). The agent also mimics the user's input line with variations, such as slightly modifying the coefficients of equations that describe the user's input line. For example, the height of a curve might change, or the angle at which a corner is made might change.

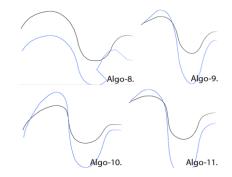At high levels of creativity (slider is set to 66-100), the

Figure 5: The drawing results from the algorithm 8-11 for high-creativity level (black lines: human, blue lines: agent)

system increases the novelty of the lines it produces using a few different techniques. For example, the system employs a line mutation algorithms that takes the input line and splices a portion of that line with a portion of another line on the canvas to introduce new content into the line (algo-8 of Figure 5). Another algorithm fits a polynomial function to part of the line that can be approximated using a polynomial (i.e. it passes the vertical line test), and then tweaks some coefficients similar to the medium creativity algorithm, but to a more drastic extent, thereby reducing the visual similarity to the input line. Yet another algorithm segments the input line into many different equations (such that it need not pass the vertical line test), and then tweaks the coefficients on some of those line segments while still maintaining the overall shape. For example, if the user drew a square, the system could produce a rectangle.

As shown in Figure 1, the user is provided with up/down voting buttons to give feedback to the system. This feedback informs the system about which algorithms in particular (within each category of creativity, i.e. low, medium, and high) the user prefers. The voting buttons are designed to train the neural network to learn the circumstances under which each type of transformation algorithm should be used. The decision about how to respond to the user's input line is a mix of the neural net analysis that compares the current input to previous responses, the creativity value, and the feedback the user has provided on each of the algorithms previously.

DRAWING APPRENTICE USER EXPERIENCE

Instead of hard-coding knowledge into the agent to make it creative, the Drawing Apprentice is designed to extract critical information from a user that is already exhibiting intelligence and creativity given the target domain in an open-ended creative task. Therefore, the most general design principle is *whenever possible, offload any higher-level cognitive tasks to the user*, i.e. enable the user to manually specify the boundaries within which the agent should operate throughout the interaction. The agent does not need to know why those parameters exist if those values help facilitate effective collaboration. Known as the *Eliza effect*, researchers have demonstrated that users attribute intentionality to virtual agents if they appear to understand the context of the situation, even when the system may not 'understand' the actions it performs [36,37]. Through interaction and user experience design, defining constraints and parameters can be smoothly integrated into their creative flow or removed altogether as the machine learning algorithms grow in complexity and sophistication.

Turn Taking

Turn taking was designed to facilitate emergent interaction dynamics, meaning the number and length of the agent's lines are dependent upon the user's recent contributions. As soon as the user ends their current line, a timer begins. If this timer passes the arbitrary value of 2 seconds before the user begins their next line, their turn ends, and the system

starts to draw. The system's turn will be approximately the same number of lines as the user's turn to mirror the interaction. However, the user may begin drawing at any point, which can lead to synchronous collaboration.

User Feedback

While humans are able to leverage extremely subtle and implicit cues during joint activity and coordination, enabling basic feedback mechanisms, such as voting, begins to approximate a similar process of coordination as human collaboration. From the perspective of optimizing the machine learning algorithms, the human should be required to provide feedback on every contribution made by the system. To maintain the flow of the creative experience, though, the human is never required to provide feedback, but rather users are given the choice to vote whenever it occurs to them, which can be highly variable between users. This presents an interesting opportunity where improving the user experience design might potentially improve the performance of the machine learning algorithms (since more feedback helps train the system).

Character Design

To simulate the dynamism and embodied nature of real-time human collaboration, the Drawing Apprentice character draws lines dynamically, meaning lines do not appear at once in full, but are gradually animated through until their completion. Dynamic line drawing is meant to provide a sense that the system is going through the embodied act of creating a line.

Pilot studies revealed the importance of having a character represent the Drawing Apprentice on the canvas, which is a sentiment echoed in the literature on embodied virtual agents [1,16]. In an early version of the prototype, the character only appeared while the system was drawing. In the pilot studies and during demos watching and talking about the character seemed to excite the participants. Multiple users requested a permanent presence, or 'home' for the character while it was not drawing. We improved the character design and created a home base for the agent to return to after it has completed its drawing. In our context, character design may be able to improve the perceived performance of the agent. With the right animations and visual design, the system might appear more creative and intelligent without any change to the algorithms.

USER STUDY

We hypothesize that a system that mimics and builds on the contributions of users in a real-time drawing collaboration can impact the sense-making process of art creation with similar benefits as participatory sense-making exhibited in open-ended creative collaboration between two people. This user study was designed to help understand participatory sense-making in the domain of collaborative drawing and delineate critical mechanisms that foster it. We investigate to what extent users can effectively work with the Drawing Apprentice in a way that enables the user to
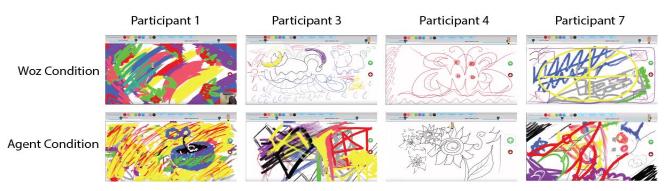
Figure 6: Representative sample of drawings from the user study. Top row: WoZ Condition; Bottom Row: Agent Condition

interactively and co-creatively build artistic meaning as the artwork develops.

Study Design

For this study, we had 7 participants, 4 female, and 3 male with an average age of 25 (ranging from 20-45) recruited from the student population at Georgia Tech. The artistic experience of the group was generally categorized as novice, with an average of 2.15 on a 5 point scale ranging from no artistic experience to 5 years of professional practice in the field. The data generated from the study included video recordings (see Figure 6), the transcribed audio data from the retrospective protocol analysis, the log data from the system, and the survey data.

The experiment was divided into two phases that each included a 12-minute collaborative drawing task, a retrospective protocol analysis, and a survey about the participant's experience interacting with the system. A non-collaborative drawing task was not included because drawing independently lacks any social coordination and participatory sense-making, which were the focus of the current study. Each experimental session lasted approximately one hour. The experiments were conducted using a Microsoft Surface tablet and a capacitive pen as input to the device. The Drawing Apprentice system was running as a web application and expanded to full screen (see Figure 1).

Participants were first oriented with the basic drawing features of the interface (line thickness, color selection, input method), as well as the unique features of the Drawing Apprentice system, such as the voting buttons and creativity slider. The experimenter described how each vote helped the system understand what the user liked, and the creativity slider controlled how 'creative' the agent was, with 0 being less creative and 100 being the most creative.

Participants were then given an open-ended prompt to 'collaborate with the system for 12 minutes to create a drawing.' One drawing task was collaborating with the Drawing Apprentice system (referred to as the agent condition), while the other task was collaborating with a Wizard of Oz agent being controlled by an expert human artist (referred to as the WoZ condition). The interface was the same in both conditions, and participants were not aware of which condition they were experiencing. They also were unaware that an expert human was controlling the system in the WoZ condition. The experimental conditions were randomly ordered to account for learning effects.

In the WoZ condition, an expert human artist controlled the system's drawing contributions from another room. The expert artist was one of the researchers from the team that has been collaborating with novices and experts in the domain of abstract drawing for over ten years. The expert used approximately the same interaction dynamics as the system, waiting 2 seconds after the user was finished drawing to begin the turn, as well as drawing approximately the same amount of lines as the user's last turn. The expert's responses were calibrated using the creativity slider, with low creativity turns closely following the participant's lines, while high creativity resulted in more novel contributions.

During the retrospective protocol analysis, the experimenter prompted the user to explain their thought process throughout the video walkthrough. The survey focused on evaluating the quality of the collaboration and how influential their collaboration was in defining artistic goals.

FindingsThis section presents the preliminary results from the retrospective protocol analysis. We performed a qualitative data analysis on the transcribed interview data using thematic analysis targeting language relating to key themes about participatory sense-making. These findings will help us develop a formal coding scheme that can be used to code the video recordings according to the elements that are identified below as important for participatory sense-making in the domain of drawing collaboration.

A key factor in this analysis is determining why contributions seem to 'make sense' to the user or 'not make sense' to the user. We expect contributions that appear completely erratic to be too far outside the meaning structures developed by the user to be integrated, while contributions that are on the fringe of the user's current meaning have the potential to expand visual ideas and guide interactions in-the-moment. Further, the manner in which interaction occurs can impact whether the agent appears to coordinate with the user, i.e. the rhythm of turn taking, the speed of lines drawn, the size of lines, etc.

Supporting evidence for participatory sense-making comes from participants reporting that the agent was able to contribute to the drawing in the following meaningful ways: (1) build on the contributions of the user, (2) demonstrate a certain degree of coordination and mutual sense-making with the user. In particular, we are interested in whether the agent can help shape these meaning structures such that *creatively engaging*, *unexpected,* and *surprising* meaning emerges through interactions over time.

Wizard of Oz Condition

The first step in the analysis was to delineate how participants described participatory sense-making in the case of human collaboration (WoZ experimental condition). In this circumstance, the thematic analysis focused on concepts related to a structural coupling between the collaborators, during which there was a mutual co-regulation of the activity, meaning both the human and the system defined basic units of meaning that were extended and added on by each participant in a way that would not have occurred without the collaboration.

Making Sense of the Agent

P2 described how the agent influenced the creative process, reporting

> "*I'm definitely taking into account what the agent is doing and some combination of trying to figure out how to sort of control the agent, and also work with it, because it did things that I wouldn't necessarily expect, which is cool.*"

P5 describes how echoing and mirroring actions helped anticipate and make sense of the agent, stating

> "*The system tended to echo what I did, and mirror the patterns I made, sometimes in different locations on the screen, but I was able to pretty quickly anticipate the kind of move, not exactly what it would do, but to the point where I was pleased with some of them, or displeased....*"

P5 elaborated this sentiment by describing he was able to find a 'comfort zone' in which the agent would make predictable contributions, "*At that point, I liked what I expected it to do. That didn't really surprise me there, but in a comforting kind of way.*" While unexpected contributions may have inspired the user, achieving basic structural couplings enabled a 'comfort zone' to emerge that provided some stability and predictability in the collaboration that helped contribute to the feeling of a 'dialog' with the system.

Interaction Dynamics

P3 described the experience of WoZ collaboration in distinct categories that help shed light on how the interaction was conceptualized:

> "*There was like me teaching it, or it copies me, or he comes up with something cool, and I copy it, or we were like collaborating on something we both know what's going on. I guess there is just another thing that is random, like I don't know what to draw, and I'll do something random, or he doesn't know what to follow and he does something random.*"

As P3 worked to make sense of the agent's behavior, they noted distinct modes of interaction that included both the agent and user copying, working on a joint activity, as well as injecting novel or 'random' contributions to help move the collaboration forward while the user or system were unsure how to contribute. Examples of joint activities include coloring (i.e. filling a shape with one color) as well as one of the collaborators (agent or user) completing each other's thoughts. Achieving joint activity relied heavily on mutual spatial awareness, which is detailed later in the Sense-Making Evaluation section.

The notion of participants copying the agent was introduced again with a strongly positive association from P4. There was a sentiment that when the participant chooses to copy the agent, it signified the agent's contribution were accepted and integrated into the user's current artistic intention, i.e. structural coupling. P4 described this sentiment,

> "*It was really cool. It is taking the accent points from the lines and making them more accented. So I was like OK, I'm going to copy you….I think throughout this interaction, there are times when the system lead me instead of the other way around... I felt a little like the system. I was like: am I just copying it?*"

The system and user were taking turns leading the interaction and suggesting new content and activities in which to engage, demonstrating a mutual co-regulation in the participatory sense-making process. Once users made sense of how the system interacts, the users engaged in a playful process of exploring the boundaries and challenging the system with increasingly 'difficult' inputs. For example, P4 'challenged' the agent, stating "*So, I was like OH! It looks like a butterfly, and I was like haha system! What are you going to do now? So I draw in moons, and I was like OK, what are you going to do?*" Instead of worrying about the outcome of the artwork, P4 was creatively engaged by the dialog that emerged with the system. This interaction inspires the participant to generate lines that test the limits of the system.

Emergent Meaning

P1 described the process of discovering emergent meaning as a result of working with the agent, saying

> "*I drew that, and it started out as random stuff, and it reminded me of a flower or a star, and then the agent drew that, and it reminded me of a halbird….so I extended it, and then the agent extended it further...*"

In this instance, the collaborators formed a brief structural coupling during which the user and agent were both contributing ideas and building onto a core idea that emerged through interaction. P6 also reports meaning emerging through a process of co-regulated structural coupling when he describes redefining his goals about the agent's contributions,

"I wanted to make a wagon, a cart thing. But then it made something like this, then I made it a bed instead. Then, it did this really nice hatching on the pillow, which I liked."

Here, P6 originally intended to draw a cart. However, based on the system's contribution, he redefined that goal as a bed, which was further elaborated by the system.

Agent Collaboration Condition

Next, we describe some illustrative examples during which participants reported participatory sense-making during collaboration with the co-creative agent. Then, we will list some of the most critical evaluation metrics users reported for determining whether or not the agent's contributions 'make sense.'

Making Sense of the Agent

P4 described how she developed a strategy for anticipating the agent's response through experimentation,

"I decided to make a spiral, and see if the AI would continue my spiral, but then it made it elsewhere, so I was like that's cool...so I drew a different spiral, to see what it would do, and they did the same thing….so I was like maybe I can use that and replicate smaller things, so I would make a flower, so I anticipated what it did."

This quote shows evidence that P4 began to make sense of the agent's reactions. P5 also discovered types of actions that elicited good responses, saying *"I did a lot of these Bezier curves because I liked the way it tried to answer me with those...."* This type of predictability was critical for the sense-making process in artistic collaboration.

Interaction Dynamics

Users were able to couple their behavior with the system, but as P5 describes, this required them to submit to how the system worked in a one-sided manner,

"This is when I sort of discovered how to work with it on that pattern over there… I'm sort of doing what it's doing, and we're feeding off of each other….I quit fighting it and started collaborating...I'm playing more by its rules. I'm sort of anticipating what I thought it would do."

This quote demonstrates participatory sense-making, but the process is being regulated mostly by the agent because the participant has to work from the agent's contribution in order to achieve 'collaboration.' This is contrasted to the mutually co-regulated process during which each partner takes turns leading and defining new goals, as seen in human collaboration. P5 surrendered a certain degree of autonomy when he decided to '*play by the agent's rules.*'

Some participants were comfortable surrendering this autonomy. However, in some cases, such as P7, dealing with the agent was frustrating, since there was "*a lot of prodding him into doing things. It was more about prodding him, and being concerned about him rather than me doing something creative.*" This type of interaction dynamic could be characterized as trying to lead or control the agent as opposed to engaging in a mutual co-regulation where each party exhibits autonomy. While prodding the agent may reduce the artistic autonomy of users, the dialogical interaction still had the capacity to propel the interaction forward, as P5 describes:

"I think if you had just given me a blank canvas and told me to draw anything, as a non-artist, I would have quickly given up, because I didn't know what I was doing. The fact that there was a back and forth, and new things were emerging on the screen, made me want to try to answer it, or to try to prod it to coming up with another addition to the drawing. I was probably more engaged than I would be if I was just drawing on a blank canvas by myself."

Emergent Meaning

P1 summarized the effect that collaboration with the co-creative agent had on their creative process, saying "*It made it less structured. I did more doodling to see how I could incorporate what the pencil was doing and see how I could interact with it.*" The 'doodling' that P1 described leads to emergent and unpredictable visual elements in the drawing. This type of interaction dynamic helped users discover novel visual ideas. P6 describes how working with the co-creative agent challenged them to move beyond their 'standard' drawing,

"It felt nice when I had to change my standard drawing, because that's something I always draw. This was more of a challenge for me to say: let's see what I can come up with to draw."

Adapting to and incorporating the agent's contributions helped participants draw in ways they would not have done without the collaboration. The challenge for the Drawing Apprentice is determining how to make better unexpected choices that encourage users to incorporate the agent's contributions rather than ignore them. Utilizing user feedback is one method to help facilitate this type of coordination.

Voting Buttons and Creativity Slider

Participants reported being able to make sense of the agent and collaborate with it in both conditions, but the voting buttons and creativity slider were not fully leveraged by

	Making Sense of Agent	**Interaction Dynamics**	**Emergent Meaning**
WoZ Collaboration	User anticipated and predicted agent responses	Co-regulated coupling where leader varies between agent and user	Visual ideas emerge from both user and agent and transform over time
Agent Collaboration	User anticipated and predicted agent responses	User leads and controls agent	User incorporates and adapts to agent responses

Table 1: Comparing participatory sense-making between the two collaboration conditions

most participants. The down vote mechanism was used to discourage the system from behaving in a particular manner, such as the style of its lines. For example, P2 reports:

> "*I definitely up voted a lot in the first time, and I down voted a lot more this time. I wanted to try to get rid of the fast lines I couldn't see, which kind of happened….I tried to discourage when it did really really shaky stuff. It kind of helped.*"

Other participants reported being uncertain about how voting affected the system. For example, when directly asked how voting affected the agent, P1 states, "*I don't know. I don't know if they really affected it.*" Other participants experimented to a further degree, but still described uncertainty when voting on the agent's contributions, with P4 reporting "*I'm not sure if the voting buttons did anything, because I down voted it, but it continued to do what it did.*" In general, participants expected immediate feedback from a down vote, such as the system immediately stopping its current drawing task. P2 elaborated one reason why voting was confusing, stating:

> "*It definitely takes several iterations of a down vote for it to figure out exactly what you are trying to discourage, like the line placement, or what type of thing you are trying to discourage, because there probably several things it considers when trying to place a line.*"

This comment highlights the fact that each vote could relate to several parameters of the line, such as its placement, the style (shaky, smooth, curvy, etc.), and the shape. Participants were uncertain how their vote would affect the system due to this ambiguity. For example, P3 stated that each system input "*has a lot of attributes, I don't know if it discards or has a decreased value of all of those, or I just don't like this particular thing.*" Providing users with explicit visual feedback about how their vote affected the algorithms would help address this shortcoming of the system.

Some participants described voting as distracting, or similarly described how they forgot about voting when they were deeply immersed in a drawing task. For example, P5 described his experience with voting, saying "*I tried to use those [voting buttons] this time to reinforce things I liked, but… I found those more distracting.*" When users were deeply engaged in collaboration, they reported thinking less about the voting and creativity mechanisms. For example, P4 stated "*I totally forgot about the creativity slider, and the voting buttons for a while because I was so intrigued by what the system was doing on its own at that level of creativity.*" These mixed sentiments about the feedback mechanisms provide insights for updating the implementation of voting and feedback that will be outlined in the Design Recommendations section.

Sense-Making Evaluation Metrics

Achieving human-level collaboration is the ultimate goal for the type of co-creative system we designed. To do so, it is critical to understand the nature of participatory sense-making in the domain of artistic collaboration and the metrics users employ to evaluate their partner. Next, we describe key concepts that emerged in the thematic analysis when we focused on contributions that appeared random or erratic in a way that was not beneficial and did not lead to emergent meaning. This dimension of the analysis will help inform the continued development of our system as well as other co-creative agents in the artistic domain. The concepts of *spatial awareness*, *visual similarity*, and *perceptual logic* emerged as particularly important when participants described how they evaluated whether contributions from the system made sense.

Spatial Awareness

One of the most common metrics users employed to evaluate whether the system's contributions made sense related to whether the agent exhibited *spatial awareness* of things that were previously drawn. According to participants, the agent should maintain an awareness of visual elements that have been created previously to not 'mess up' what the users have drawn. Participants became frustrated when they perceived the agent as 'messing up' their contribution rather than building off of it. For example, P1 states: "*Usually if I draw actual things, they keep within lines, and you know, you organize the space, and when you have other things encroaching on it, it's hard.*" P4 reported a similar occurrence while collaborating with the agent, stating: "*The system would draw on top of other things we drew, when I specifically avoided those things…so that was annoying.*" It was obvious that the system did not understand what the user's intention was when it drew on top of the participant's lines.

Visual Similarity

The agent's contribution should retain some visual similarity of the user's contribution so that the user can help understand the relationship between their contribution and the system's. When the user was able to see the visual similarity between the agent's and their behaviors, they were able to reason through why it made a contribution, which can reduce their frustration if it 'messes up' something they drew. For example, P2 states: "*This is cool. I like this pattern here. I did that, and he did a more acute angle.*" The agent's contributions that were on top of the user's lines were more acceptable if the user could justify why the agent had created those lines. Typically, visual similarity and close spatial proximity to the user's input lines helped provide clues to participants about why the agent drew a particular object. For example, P3 states "*I was trying to draw two birds, but he just tries to copy…I didn't really like it, because it's not really a bird anymore, but I get why he did that.*" It is important to understand the context in which a contribution is accepted even though the participants report not 'liking' it. This acceptance can lead to emergent meaning as users work to continue to transform the content rather than completely abandoning or ignoring

the contribution, as is sometimes the case when the user perceives contributions to be too erratic.

Perceptual Logic

The agent's lines are evaluated as 'messing up' the participants drawing when the agent drew over top of a previous structure without taking into account any of the features of that structure. When the agent did make contributions within or on top of a previously defined shape, it must adhere to what could be termed the prevailing 'perceptual logic' of that region [8,9]. Perceptual logic is a concept that describes how each region or idea in a drawing has its internal set of rules and mechanism that serve to constrain what type of visual contributions seem relevant and logical for the target region. It could be conceptualized as the local style of a particular region.

In some cases, perceptual logic appeared to be based purely on visual relationships, like the repetition of similar patterns (as noted in the previous section on visual similarity), while in other cases, it was based on a semantic definition of the object, such as a house that has different components that are 'right' or 'wrong.' P3 described a perceptual logic of 'orderly geometric shapes', reporting:

> "*I didn't like too much what was going on there because I was trying to actually do like geometric shapes and stuff, so I was thinking something more orderly, so I just moved back and started drawing regular shapes.*"

Contributions that did not adhere to the perceptual logic of 'orderly geometric shapes' did not make sense to the participant. Another example of perceptual logic is when P5 reported:

> "*I didn't like the way it looked. It messed up my swirl, and it wasn't symmetric....if it was going to do it, it should be centered in the spiral, then I might have thought it would have been ok...it seemed arbitrary.*"

The negative sentiment this participant had for that particular violation shows that sometimes perceptual logic is tied up with an emotional response, evidenced when P5 reiterates his sentiments on the spiral, saying "*I really hated that thing inside the spiral, that's what I hated the most.*"

DESIGN RECOMMENDATIONS

Given its prominence in evaluating the agent's behavior, spatial awareness seems to be a foundational skill for a co-creative drawing partner. One method for achieving spatial awareness is by constraining the agent's learning mechanisms to particular regions of the artwork, i.e. executing certain types of drawing behaviors in certain regions based on user feedback. Since drawings develop over time, the perceptual logic in each region is also subject to change as regions grow, transform, and potentially connect with other regions. Given this dynamism, the system's learning algorithms should be temporally sensitive as well. For example, when users re-visit previously established regions, they may be trying to accomplish much different tasks, which could drastically change the

perceptual logic that is appropriate. Considering both space and time in the learning algorithms should therefore improve the agent's ability to coordinate with users and engage in participatory sense-making.

In both conditions, participants used repetition to try to train and reinforce behaviors in the agent, but this type of implicit feedback was not registered by the co-creative agent. By analyzing the relationship of the user's input lines for a given turn, it might be possible to classify what type of contribution the user is making, i.e. defining a pattern through repetition, drawing a complete shape, or beginning a joint activity, such as coloring. All these modes of teaching emerged during the study, but the agent's machine learning architecture was not specifically designed to learn in that manner. Additionally, other implicit cues observed during interaction could serve as feedback for the system, such as when the user copied the agent's contribution. Co-creative agents should take this type of mimicry as positive feedback indicating the early stages of a structural coupling that can facilitate participatory sense-making.

Our findings indicated that users did not fully understand how using voting and the creativity slider affected the behavior of the system. This could be mitigated by providing more explicit feedback about how their vote affected the agent's knowledge and drawing behaviors. In addition to describing how votes affect the creativity of the system, our findings indicated that users would benefit by reducing the ambiguity of binary feedback. Disambiguating user feedback could include providing a more continuous evaluation scale (versus the current binary like/dislike), as well as categories of feedback, such as providing independent feedback on the location, style, and content of the agent's drawing contribution.

CONCLUSIONS

This paper described a co-creative drawing partner called the Drawing Apprentice. The system was designed to improvise and collaborate with users in real time on an abstract drawing. We presented the conceptual framework of participatory sense-making from the cognitive science theory of enaction to help understand and categorize the interaction dynamics involved in collaboration. That conceptual framework was then used to compare human-human collaboration and human-computer collaboration during a user study with the Drawing Apprentice system. Our analysis identified how meaning emerged through participatory sense-making in both conditions. We identified critical metrics users employ to evaluate whether the system's contributions 'made sense' and describe how participants worked to train and provide feedback to the system to help coordinate their interactions. We leveraged these findings to propose design recommendations for co-creative agents.

REFERENCES

1. Baylor, A. L. Promoting motivation with virtual agents and avatars: role of visual presence and appearance. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 364, (2009) 1535: 3559–3565.

2. Biles, A. J. GenJam: A genetic algorithm for generating jazz solos. *Proceedings of the International Computer Music Conference*, Citeseer, (1994), 131–131.

3. Boden, M. A. *The Creative Mind: Myths and Mechanisms*. Weidenfeld & Nicolson, London. 1990.

4. Charman, T., Swettenham, J., Baron-Cohen, S., Cox, A., Baird, G. and Drew, A. Infants with autism: an investigation of empathy, pretend play, joint attention, and imitation. *Developmental psychology* 33, 5 (1997), 781.

5. Colton, S., Wiggins, G. A. Computational creativity: The final frontier? *ECAI*, (2012) 21–26.

6. Compton, K. and Mateas, M. Casual Creators. *Proceedings of the Sixth International Conference on Computational Creativity June*, 228. 2015.

7. Davis, N., Comerford, M., Jacob, M. Hsiao, C.-P. and Magerko, B. An Enactive Characterization of Pretend Play. *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, (2015), 275–284.

8. Davis, N., Do, E. Y.-L. Gupta, P. and Gupta, S. Computing harmony with PerLogicArt: perceptual logic inspired collaborative art. *Proceedings of the 8th ACM conference on Creativity and cognition*, (2011) 185–194.

9. Davis, N., Hsiao, C.-P., Popova, Y. and Magerko, B. An Enactive Model of Creativity for Computational Collaboration and Co-creation. In *Creativity in the Digital Age*. Springer, (2015), 109–133.

10. Davis, N., Hsiao, C.-P., Singh, K. Y., Li, L, Moningi, S. and Magerko, B. Drawing Apprentice: An Enactive Co-Creative Agent for Artistic Collaboration. *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, (2015), 185–186.

11. Dixon, D., Prasad, M. and Hammond, T. iCanDraw: using sketch recognition and corrective feedback to assist a user in drawing human faces. *Proceedings of the 28th international conference on Human factors in computing systems*, (2010), 897–906.

12. Edmonds, E. and Candy, L. Computer support for creativity. *International Journal of Human-Computer Studies* 63, (2005), 4-5:

13. Fantasia, V., De Jaegher, H. and Fasulo, A. We can work it out: an enactive look at cooperation. *Frontiers in psychology* 5. 2014.

14. Flagg, M. and Rehg, J. M. Projector-guided painting. *Proceedings of the 19th annual ACM symposium on*

User interface software and technology, ACM Press, (2006), 235–244.

15. Fuchs, T. and De Jaegher, H. Enactive intersubjectivity: Participatory sense-making and mutual incorporation. *Phenomenology and the Cognitive Sciences* 8, 4 (2009), 465–486.

16. Gulz, A. and Haake, M. Design of animated pedagogical agents—A look at their look. *International Journal of Human-Computer Studies* 64, 4 (2006), 322–339.

17. Hodhod, R., Piplica, A. and Magerko, B. A formal model of shared mental models for computational improvisational agents. *Proceedings of the 12th Annual Conference on Intelligent Virtual Agents*. 2012.

18. Hodhod, R., Magerko, B. and Gawish, M. Pharaoh: Context-Based Structural Retrieval of Cognitive Scripts. *International Journal of Information Retrieval Research (IJIRR)*, 2, 3 (2012), 58–71.

19. Hodhod, R. and Magerko, B. Pharaoh: Conceptual Blending of Cognitive Scripts for Computationally Creative Agents. *The Twenty-Seventh International Flairs Conference*. 2014.

20. Hodhod, R. and Magerko, B. Closing the Cognitive Gap between Humans and Interactive Narrative Agents Using Shared Mental Models. *21st Annual Intelligent User Interfaces Conference*. 2016.

21. Hoffman, G. and Weinberg, G. Shimon: an interactive improvisational robotic marimba player. *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, (2010), 3097–3102.

22. Jacob, M., Coisne, G., Gupta, A, Sysoev, I., Gav Verma, G., and Magerko, B. Viewpoints AI. *AIIDE*. 2013.

23. Jacob, M., Zook, A. and Magerko, B. Viewpoints AI: Procedurally Representing and Reasoning about Gestures. *Proceedings of DiGRA*. 2013.

24. De Jaegher, H. and Di Paolo, E. Participatory sense-making. *Phenomenology and the cognitive sciences*, 6, 4 (2007), 485–507.

25. Lee, Y. J., Zitnick, C. L. and Cohen, M. F. ShadowDraw: real-time user guidance for freehand drawing. *ACM Transactions on Graphics (TOG)*, 27. 2011.

26. Lubart, T. How can computers be partners in the creative process: Classification and commentary on the Special Issue. *International Journal of Human-Computer Studies* 63, 4-5 (2005), 365–369.

27. Magerko, B., Dohogne, P and DeLeon, C. Employing Fuzzy Concept for Digital Improvisational Theatre. *AIIDE*. 2011.

28. O'Neill, B., Piplica, A., Fuller, D. and Magerko, B. A Knowledge-Based Framework for the Collaborative Improvisation of Scene Introductions. In *Interactive Storytelling SE - 10*, Mei Si, David Thue, Elisabeth

André, JamesC. Lester, Joshua Tanenbaum and Veronica Zammitto (eds.). Springer Berlin Heidelberg, (2011), 85–96.

29. Piplica, A., DeLeon, C. and Magerko, B. Full-body gesture interaction with improvisational narrative agents. *Intelligent Virtual Agents*, (2012), 514–516.

30. Sawyer, R. K. and DeZutter, S. Distributed creativity: How collective creations emerge from collaboration. *Psychology of Aesthetics, Creativity, and the Arts* 3, 2 (2009), 81.

31. Schertz, H. H. and Odom, S. L. Promoting joint attention in toddlers with autism: A parent-mediated developmental model. *Journal of autism and developmental disorders* 37, 8 (2007), 1562–1575.

32. Schneiderman, B. Accelerating Discovery and Innovation. *Communications of the ACM* 50, 12. 2007.

33. Shneiderman, B., Fischer, G., Czerwinski, M. et al. Creativity Support Tools: Report From a U.S. National Science Foundation Sponsored Workshop. *International Journal of Human-Computer Interaction*, 20, 2 (2006), 61–77.

34. Stewart, J. R., Gapenne, O. and Di Paolo, E. *Enaction: Toward a new paradigm for cognitive science*. MIT Press. 2010.

35. Voigt, M., Niehaves, B. and Becker, J. Towards a Unified Design Theory for Creativity Support Systems. GroupSystems (2011), 152–173.

36. Wardrip-Fruin, N. *Expressive Processing: Digital fictions, computer games, and software studies*. MIT press. 2009.

37. Weizenbaum, J. ELIZA---a computer program for the study of natural language communication between man and machine. *Communications of the ACM 9*, (1966), 36–45.

38. Wiggins, G. A. Searching for computational creativity. *New Generation Computing* 24, 3, (2006), 209–222.

39. Wiggins, G. A. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19, 7 (2006), 449–458.