# Building Artistic Computer Colleagues with an Enactive Model of Creativity


**Nicholas Davis[1], Yanna Popova[2], Ivan Sysoev[1],**

**Chih-Pin Hsiao[3], Dingtian Zhang[1], Brian Magerko[1]**

[1]School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA USA
{ndavis35, ivan.sysoev, alandtzhang, magerko}@gatech.edu

[2]Department of Cognitive Science
Case Western Reserve University
Cleveland, OH USA
yanna.popova@case.edu

[3]College of Architecture
Georgia Institute of Technology
Atlanta, GA USA
chsiao9@gatech.edu



## Abstract

This paper reports on the theory, design, and implementation of an artistic computer colleague that improvises and collaborates with human users in real-time. Our system, Drawing Apprentice, is based on existing theories of art, creative cognition, and collaboration synthesized into an enactive model of creativity. The implementation details of the Drawing Apprentice are provided along with early collaborative artwork created with the system. We present the enactive model of creativity as a potential theoretical framework for designing creative systems involving continuous improvisational collaboration between a human and computer.


## Introduction

Creative technologies have come a long way in supporting human creativity in a variety of ways. Modern creativity support tools (CST) have been extremely effective at helping users produce higher quality products by allowing them to explore creative possibilities, perform complex simulations, and record and track ideas (Shneiderman 2007). However, with all their capabilities and features, popular creativity support tools like Adobe's Photoshop are not yet able to generate original artistic contributions, such as new lines or brush strokes that add to the user's artwork. Recent advances in artificial intelligence and computational creativity are beginning to change this by developing co-creative computer colleagues to enrich the human creative process in a completely new manner through collaboration with a creative computer (Lubart 2005).

Computer colleagues can bridge the gap between CSTs that support a creative person and computers that generate creative products autonomously (see Figure 1). We hypothesize collaboration with computer colleagues based on the enactive model of creativity can enrich the creative process like human collaboration (i.e. increase playful exploration, motivation, creative engagement) in open-ended creative domains such as non-representational visual art. We have designed and implemented a prototype of an artistic computer colleague using the enactive model of creativity (EMC) to test this hypothesis.

Our system, called Drawing Apprentice applies EMC to abstract improvisational art. This artistic domain was selected for its open-ended, flexible and emergent art patterns (Clouzot 1956). EMC synthesizes several cognitive science and creativity theories to model creativity as an enactive process that emerges through constant interaction with the environment and other agents within it. In this view, creative actions emerge through experimental interactions with the environment based on simulations and perceived artistic affordances rather than executing a fully formed plan and artistic goal.

In the following sections, we first introduce co-creativity in the context of computational creativity and improvisational abstract art. Next, we provide some background on enactive cognition. Then, we present our enactive model of creativity and show how it helped us design an improvisational drawing agent. Finally, we consider evaluation metrics and show early artwork created with the system.

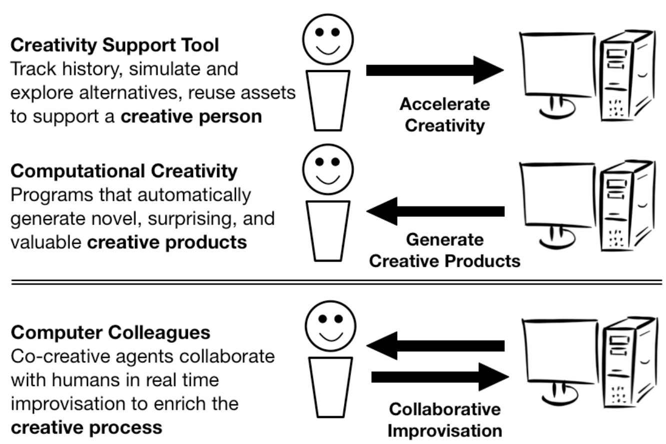


Figure 1: Computer Colleagues Bridge the Gap Between CSTs and Computational Creativity

# Background

## Computational Creativity

HCI researchers build creativity support tools that augment and extend the creative abilities of humans (Shneiderman 2007), while AI researchers develop computationally creative systems that implement and sometimes elaborate on cognitive theories of creativity (Boden 2003; Colton 2008; Li et al. 2012). Enormous progress has been made in these two complementary pursuits; however, there is a gap in the research literature about blending humans and computers in a continuous and collaborative co-creative process (Lubart 2005). The field of computational creativity does not yet have a guiding paradigm or set of design principles to structure creative systems involving continuous real time improvisational collaboration between creative humans and creative agents (Lubart 2005).

Co-creativity is classified as multiple parties contributing to the creative process in a blended manner (Candy et al. 2002). It arises through collaboration where each contributor plays an equal role. Cooperation, on the other hand can be modeled as a distribution of labor where the result only represents the sum of each individual contribution (Candy et al. 2002). Co-creativity allows participants to improvise based on decisions of their peers. Ideas can be fused, and built upon in ways that stem from the unique mix of personalities and motivations of the team members (Candy et al. 2002). Here, the creative product emerges through interaction and negotiation between multiple parties, and the sum is greater than individual contribution. These interaction principles can be extended to include a sufficiently creative agent that can collaborate with human users in a new kind of human-computer creativity.

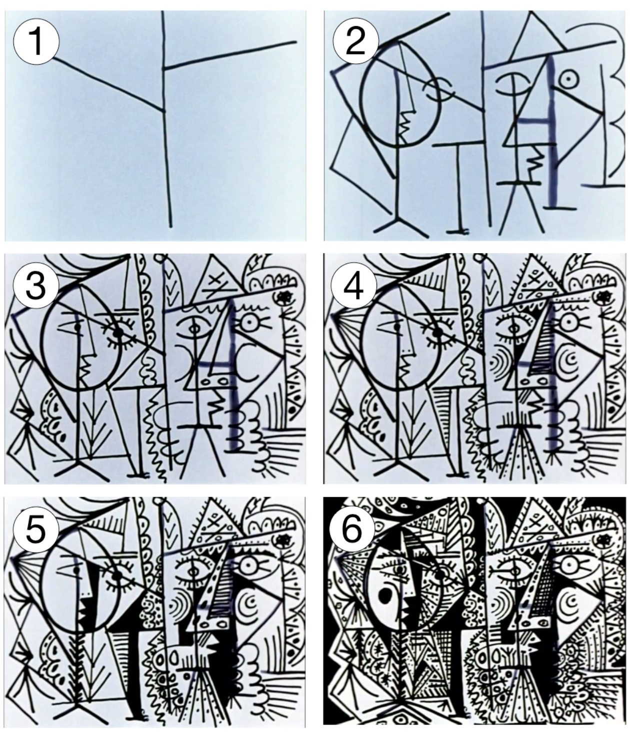


Figure 2: Time-lapse representation of Picasso's abstract art improvisation creative process reproduced from a film of Picasso painting in Clouzot (1956)

Some approaches that have yielded interesting examples of human-computer creativity include mimicry, structured improvisation, and using contextual clues to negotiate shared mental models. The improvisational percussion robot Shimon mimics human musicians by analyzing the rhythm and pitch of musical performances and generating synchronized melodic improvisations (Hoffman & Weinberg 2010). In practice, the human and robot develop a call-and-response interaction where each party modifies and builds on the previous contribution. Some co-creative agents use sensory input to construct mental models of agents, actions, intentions, and objects in the environment (Magerko et al. 2010). Mental models help agents effectively structure, organize, interpret, and act on sensory data in real time, which is critical for meaningful improvisation.

## Abstract Improvisational Creativity

Pablo Picasso's work is the most well known example of the type of improvisational abstract art the system was designed for. One of the defining features of abstract art is its ability to morph and transform throughout the creative process as the artist discovers, assigns, and re-interprets meaning in the artwork (see Figure 2 and Clouzot's (1956) *Le mystère Picasso* for additional context).

In the cognitive science literature, this type of meaning re-assignment is referred to as a *conceptual shift* (Nersessian 2008). Colloquially termed the Eureka! or Aha! moment, conceptual shifts occur when two separate knowledge domains are connected in the mind (Boden 2003; Nersessian 2008). It is often partially or wholly responsible for insights that lead to creative discoveries and solutions.

Abstract art is particularly interesting for creativity research because conceptual shifts and flexible meanings are its cornerstones. Its fluidity makes abstract art ideal for collaboration, as collaborators quickly and easily negotiate common ground and construct shared meaning in an artwork. Abstract art contributions also cannot be 'wrong' in the same strict sense as representational art because accurate representations are not the goal, which helps lower the barrier of entry for novices (both human and computer).

Improvisational creativity more closely resembles a dialogue where each party makes contributions that feed into an interactive creative process (Sawyer 2012). Jazz improvisation exemplifies artists working together to experimentally negotiate creative strategies based on current musical themes, patterns, and the history of interaction (Sawyer 2012; Mendonça 2004).

Improvisational creativity is distinguished from other types of creativity because the product is usually ephemeral—*the process is the product* (Sawyer 2012). Computer colleagues can enrich the *creative process* by engaging artists in a fun and interesting collaborative art making experience. The final creative product could be thought of as merely a record of that collaborative experience.

## Enactive Cognition

Enactive cognition is an outgrowth of the embodiment paradigm in cognitive science. Embodiment claims cognition is largely structured by the manner in which our bodies enable us to interact with the environment (Varela et al 1991). This approach is contrasted with earlier cognitive theories that conceptualized the mind as a machine and cognition as a complex but disembodied manipulation of symbolic representations (Newell 1959). In particular, enaction emphasizes the role that perception plays in guiding and facilitating emergent action (De Jaegher 2009). In the following sections, we describe how the enactive approach reframes perception into an active and dynamic process critical for participatory sense-making, i.e. negotiating emergent actions and meaning in concert with the environment and other agents. Next, we examine the role of goals and planning in the enactive perspective. Finally, we review some sketching and design research to show evidence that enaction plays a key role in the creative process when creative individuals 'think by doing.'

**Enactive Perception** In the enactive view, cognition is seen as a cycle of anticipation, assimilation and adaptation, all of which are embedded in and contributing to a continuous process of perception and action. Perception is not a passive reception of sensory data, but rather an active process of visually reaching out into the environment to understand how objects can be manipulated (Gibson 1986; Noë 2004). This type of enactive perception minimally involves a negotiation among the following factors: 1) The subject's intentional state; 2) The skills and bodily capabilities of the individual; and 3) Perceptually available features of the environment that afford different actions such as size, shape, and weight (e.g. is it graspable, liftable, draggable, etc. as elaborated in Norman (1999)). Sensory data enters the cognitive system and irrelevant data is suppressed and filtered (Gaspar 2014). Objects and details of the environment that relate to the subject's intentional goals appear to conscious perception as affordances, which can grab, direct, and guide attention and action (Norman 1999).

Each time the individual physically moves through the environment, or acts upon the environment, that action changes the perceptually available features of the environment, which can reveal new relationships and opportunities for interaction. For example, when a painter steps back from her painting, two things happen: (1) she disengages from her current painting activity, and (2) she changes the sensory input to her visual system. From this new perspective, the artist can evaluate global relationships between local regions in the painting and discover new themes and artistic goals that can guide her next artistic decisions once she re-engages the artwork.

**Participatory Sense-Making** The enactive view accentuates the participatory nature of meaning generation, often called participatory sense making. Cognitive systems generate meaning by active transformational and not mere informational interactions with the environment (Varela et al. 1991; Gapenne and Di Paolo 2010). Each interaction with the environment can (and often does) reveal new goals, which leads to a circuitous rather than a linear creative process. Creative individuals engage in a dialogue with the materials in their environment (and other agents) to define and refine creative intentions (Schon 1992). This view is helpful in open-ended domains where goals are often discovered rather than explicitly defined.

In human daily interactions, for example, there is evidence that some form of natural coordination takes place in the shape of movement anticipation and synchronization. A good example of participatory sense-making would be the familiar situation where you encounter someone coming from the opposite direction in a narrow passageway (De Jaegher 2009). While trying to negotiate a safe and quick passage, both participants look toward their intended path (providing a social cue) while also trying to assess the projected path of other agents. Interaction then, in the form of coordination of movements, is the decisive factor in how quickly the individuals achieve their goal of passing each other.

Rather than adopting a plan with a fixed and concrete goal state to control locomotion, an enactive analysis would posit that individuals remain flexible throughout the situated action by dynamically accommodating the choice of the other agent. If the interaction cannot be settled by subtle perceptual negotiation, more intentional gestures can be recruited to communicate intention more precisely. If collision seems unavoidable, even after clear gestures to communicate intention are made, language may be recruited to settle the navigational issue with a solid plan, usually followed by a brief period of uncomfortable laughter (because we usually manage these situations without such extreme measures).

**Goals as Socially Negotiated, Dynamic, and Emergent** Even at the level of social interaction with an intelligent agent, an enactive approach tries to avoid postulating high-level cognitive mechanisms at the core of our intersubjetive skills. Enaction breaks away from traditional cognitive science theories positing precisely formulated goals, detailed planning procedures, and robust internal representations of both (Newell 1959). The co-evolution of a communicative/creative process is seen here as a gradual unfolding in real time of a dynamic system spanning a human subject, the environment, and agents within it. In this view, intentions emerge but are also transformed in and through the interaction with other agents and the environment.

One argument against a naïve planning approach in AI is that it takes a significant amount of cognitive effort to construct mental simulations that provide the level of detail and granularity required to carefully plan every complex action humans engage in (De Jaegher 2009). There is considerable evidence that demonstrates humans do, in fact, have a keen skill for visual thinking, but it still takes cognitive resources to perform mental operations and inferences on images (Kosslyn 1980). It is often simply easier to act on the environment and experiment with how different interactions affect the system (Noë 2004).

**Thinking By Doing** The literature on creativity supports the enactive perspective with research on 'thinking by doing.' There is a multitude of evidence demonstrating how both representational and non-representational artists plan their artworks using sketches, studies, and other ways to simulate artistic alternatives (Mace 2002). Sketching reduces cognitive load and facilitates perceptually based reasoning (Schön 1992). Artists generate vague ideas and then use some form of sketch or prototyping activity to creatively explore, evaluate, and refine artistic intentions (Davis 2011). Sketching allows creative individuals to think by doing. When an action or idea is materialized in some way, the perceptual system is rewarded with richer data than pure mental simulations and abstract reasoning. Additionally, cognitive resources that would have been used to simulate the action (i.e. consciously visualizing the situation) are now freed for other tasks such as interpretation and analysis (Shneiderman 2007).

## Enactive Model of Creativity

An enactive model of creativity proposes creativity as an emergent negotiation between agents with intelligent perceptual systems, exploratory interaction, and an environment rich with affordances. We first explain the visual conventions of the enactive model of creativity and describe how it can be applied to model creative cognition through time. Then, we introduce a new concept derived from our model called perceptual logic, which is a perceptual filter that highlights relevant affordances in the environment while suppressing irrelevant affordances.

## Model Description

In the enactive model of creativity (see Figure 3), the awareness of the agent is represented by the vertical rectangle situated on a spectrum of cognition, which means that the agent is 'aware' of what is perceived and its current intention. Perception is constituted partly by the mental model the agent has constructed for the current situation (top-down cognition) as well as the sensory input coming from the environment (bottom-up cognition) (Gibson 1988; Glenberg 1997; Varela et al. 1999; Stewart et al 2010; Gabora 2010).

To get a sense of the intended dynamism of this model, imagine the entire 'awareness' rectangle (the central part of Figure 3) can shift to the left or right of the cognitive continuum as a function of the agent's concentration. Routine actions only require minimal thought and a limited amount of highly relevant sensory data. The enactive (and temporally extended) model of routine actions, such as driving, would by visually depicted by having the awareness rectangle resting at equilibrium in the center of the spectrum with small deviations to the left to update and revise strategy, and deviations to the right to interactively evaluate those ideas in a perceive-act cycle (see Figure 4).

If the agent is performing an unfamiliar task, however, cognitive resources are recruited to actively build a mental model of the situation, which requires performing experimental interactions, closely examining the results in the

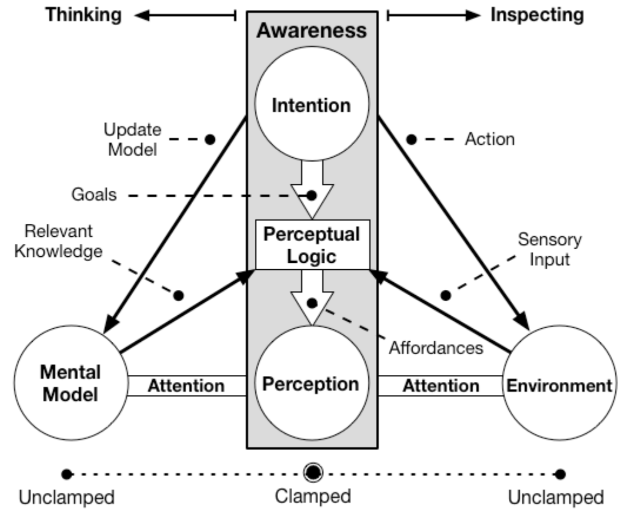


Figure 3: Enactive Model of Creativity

environment, and then updating the mental model in a slower perceive-think-act cycle. As novices learn to filter irrelevant sensory details and operate effectively with minimum conscious supervision of a task, the perceive-think-act cycle gradually tightens until expertise is achieved. Additionally, the agent can engage in pure reflection or pure interactive inspection, which would be described by tight cycles on either end of the spectrum (see Figure 4).

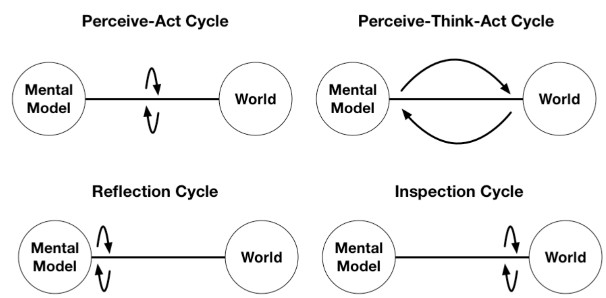


Figure 4: Cycles of cognition in the enactive model of creativity

To simulate working memory, the agent only has a limited amount of cognitive resources. These resources are used through a process of directed attention, i.e. concentration. During this simulated form of concentration, agents devote their attention to reflecting on the situation (building more detailed mental models, running complex mental simulations, etc.) and acting in a deliberate and interactive manner to inspect the world.

## Perceptual Logic

The contents of perception vary based on an individual's position on this continuum of cognition (Glenberg 1997). As individuals deviate from the equilibrium in the center of the spectrum, perception becomes partially 'unclamped,' which loosens semantic constraints on sensory input and memory (Glenberg 1997). In our model, different points on the cognitive spectrum result in a unique *perceptual logic* that is used to intelligently perceive affordances in the environment. The enactive approach in cognitive science describes the 'intelligence' of perception in a theoretical sense, but operationalizing the theory required explaining the implicit black box mechanism that makes perception 'intelligent.' The mechanism basically serves to to filter all possible affordances and present only relevant affordances to conscious perception. Perceptual logic is our proposed method for developing 'intelligent' perception in an agent.

The enactive approach proposes that perceptual intelligence arises through the formation of percept-action pairings that are chunked and internalized for quick retrieval (Noë 2004). Perceptual logic is a proposed cognitive mechanism that filters sensory data, identifies relevant percept-action pairings, and presents these percept-action pairings as affordances to perception. Perceptual logic performs a similar role as the 'simulator' in Perceptual Symbol Systems (Barsalou 1999). The simulator activates all the associated information related to a percept, including the various ways it can be interacted with based on experiential knowledge and physical characteristics.

**Clamping Perception** Research indicates that perception filters irrelevant sensory input to reduce distractions and facilitate everyday cognition (Gasper 2014). When the agent is engaged in a routine task and following well established affordances, sensory data is 'clamped' to filter out unnecessary details and un-conventional ways of seeing objects (Glenberg 1997). Everyday cognition is represented in EMC by situating the awareness rectangle in the center of the spectrum of cognition, creating a point of equilibrium. Shifting either to the left or right on this spectrum requires the agent to concentrate on either the details of her mental model or closely inspect details in the environment. At equilibrium, EMC proposes that perception is clamped to a combination of sensory input and cognitive input that optimizes routine interactions. When minor problems arise, such as small improvisational adjustments to the action based on environmental feedback, this equilibrium is slightly perturbed. The agent could generate various alternative actions by thinking (moving slightly left on the spectrum) and explore various ideas by interacting with the environment (moving slightly right on the spectrum).

**Unclamping Perception** If there is a severe disruption to the current task (e.g. a great new idea, distraction, or some kind of failure), it might become necessary to disengage from the current task to re-evaluate the situation (Dourish 2004). When an individual 'disengages' from a task, perception becomes 'unclamped' and attention shifts to thinking and simulating solutions (moving far left on spectrum) and closely examining the detail of the environment to discover new affordances (moving far right on the spectrum). The degree of concentration devoted to thinking about or acting on the environment determines how far, in either direction, awareness is situated on the spectrum of cognition. At the extreme left of the continuum (thinking) would be closing one's eyes to try to think deeply about a topic, which removes sensory input from perception altogether. At the extreme right of the continuum (inspecting) would be an individual fully concentrated on acting skillfully, carefully, and deliberately on the environment.

**Modulating Semantic Constraints** During these periods of disengaged evaluation, EMC proposes that the semantic constraints for recalling associated ideas from memory and interpreting elements in the environment become 'unclamped' to enable re-conceptualization. Unclamping semantic constraints helps overcome functional fixedness, which is a phenomenon where individuals have trouble dissociating objects from their entrenched meaning during insight problem solving (Adamson 1952).

Interestingly, this model identifies an important role for distraction in the creative process. Distraction is one way to prompt an individual to disengage from everyday cognition. In abstract art, for example, unfinished segments of the artwork (or unexpected contributions from a collaborator) may distract the artist while they are drawing. These newly discovered areas might not align with the artist's current intention. As a result, the artist might want to resolve that tension by drawing additional lines, which can catalyze the creative process. However, too many distractions might frustrate an artist.

EMC accounts for meaning negotiation by describing how perception employs different types of perceptual logic to filter affordances in the environment. Applying a different perceptual logic changes the manner in which sensory inputs are processed, organized, and made sense of. It therefore reveals different affordances in the environment, which can help the individuals discover new creative uses for objects that are relevant to goals.

## Drawing Apprentice System Design

The enactive model of creativity informs the Drawing Apprentice's cognitive architecture, and collaborative drawing and jazz improv informs the turn-taking strategies (Mendonça & Wallace 2004; Pressing 1984). Figure 5 shows the

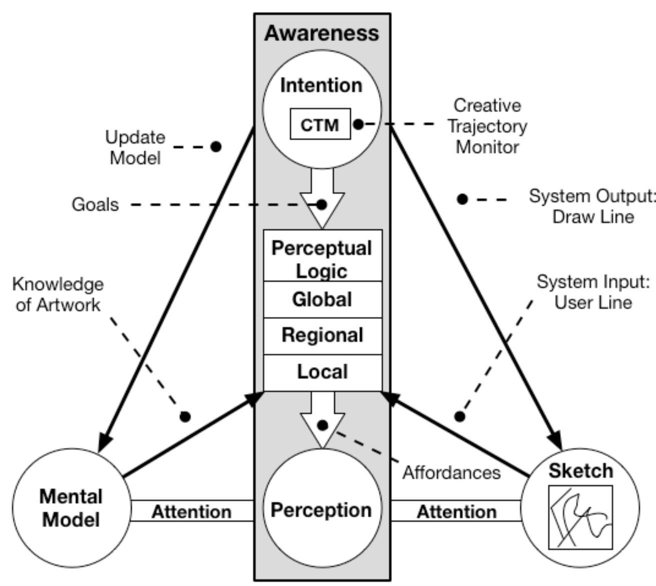


Figure 5: Apprentice Software Architecture

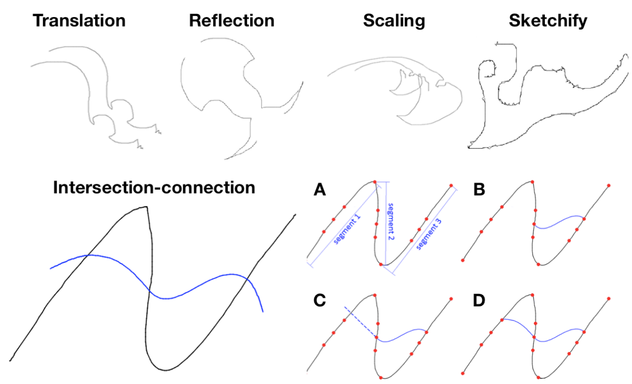


Figure 6: Local (top row) and regional (bottom row) perceptual logic drawing algorithms in system prototype.

system architecture of Apprentice. The creative dialogue begins as the human inputs a line. All current lines from the canvas are sent to the perceptual logic module. The perceptual logic module consults the creative trajectory monitor to determine what perceptual logic to apply to its current data set. The planned creative trajectory monitor has a coarse grained record of the previous drawing behavior based on the time between the user's lines (i.e. longer periods of rest represent reflection, which is categorized as global perceptual logic, and short and rapid detail strokes are categorized as local perceptual logic). The creative trajectory monitor then averages the last 10-15 seconds of user drawing behavior and selects the dominant perceptual logic of the user. The average creative trajectory is adopted by the system to determine what layer of perceptual logic to apply in the current interaction.

## Layers of Perceptual Logic

EMC suggests that each layer of perceptual logic should generate unique artistic affordances from the same input, such as shading a circle, intersecting it, and replicating it. Each logic layer sends its algorithms different amount of lines and different features for discriminating lines. There are several critical points that each perceptual logic filter can use in different ways, such as inflection points, start point, end point, segments between inflections, and corners. Moreover, gestalt groupings (e.g. proximity, similarity, closure, etc.) provide additional features to generate unique affordances building relationships between lines, groups of lines, regions, and patterns (Arnheim 2001).

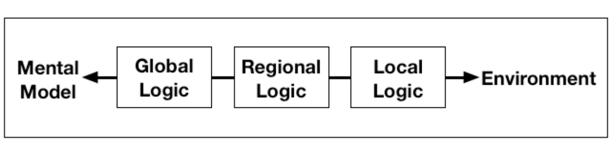


Figure 7: Layers of perceptual logic. Local perceptual logic mimics the *last input line* without any model of the artwork. Regional perceptual logic analyzes *recent input lines into* gestalt groupings to build on regional relationships. Global perceptual logic analyzes *all lines* in the agent's mental model of the artwork to evaluate overall composition and identify opportunities.

There are three layers or types of perceptual logic in EMC (local, regional, and global) determined by the position of awareness on the spectrum of cognition (see Figure 7 for an explanation of the categories of perceptual logic). We are implementing the EMC in steps with one layer of perceptual logic implemented per step. Each successive layer of perceptual logic considers a larger portion of the drawing at a higher level of conceptual abstraction (global being the most complex), which presents additional technical hurdles. Layering our implementation strategy allows a practice-based approach that encourages iterative testing with artists to ensure a meaningful artistic tool.

With only the first two layers (partly) implemented, the system can receive line input from the user, analyze it and generate an improvised response line based on the visual features of the input line and surrounding region. Table 1 and Figure 6 display the first five types of drawing algorithms we implemented in the prototype.

Local perceptual logic considers the visual features of one line. These drawing algorithms perform simple mathematical transformations on the input line and then redraw it, such as translation, reflection, scaling, and sketchify (see Figure 6). Local perceptual logic essentially mimics the creative input of the user by repeating the user's action with a small variation.

Regional perceptual logic, on the other hand, segments recent line inputs into line groups, regions, and containers based on principles of gestalt grouping, such as proximity, similarity, common fate, and continuity (Arnheim 2001). The system then generates a line that builds relationships between objects in the same region or container. Intersection-connection is the first regional algorithm that analyzes an input line into critical regions to respond to the actual shape of the line (shown in Figure 6).

Global perceptual logic (not yet implemented) considers the artwork as a whole. These algorithms are more 'intelligent' than regional and local perceptual logic algorithms because they consider how the different regions of the drawing balance to form an overall composition. When this perceptual logic is applied, the system may decide to completely decouple its contribution from the human's recent input, i.e. it can select non-active regions of the artwork on which to operate if it presents more rewarding artistic opportunities. Global perceptual logic is the highest level of cognitive functioning and will eventually include semantic knowledge such as how to draw a dog, cat, person, etc.

## System Evaluation

While creativity support tools typically help users produce a more polished product in less time, computer colleagues aim to support the *creative process* by increasing playful exploration, motivation, and creative engagement. Evaluating computer colleagues therefore involves analyzing and measuring creative engagement in the co-creative process rather than judging the creativity of the final product.

Figure 8 shows an early practice-based art study of an expert artist (the first author) collaborating with the Drawing Apprentice over a period of 2 hours. Drawings 1 & 2 demonstrate short turn taking collaboration between the artist and the Drawing Apprentice (computer lines are blue). Without the regional and global perceptual logic layers, the system only has minimal knowledge of the artwork. It knows what each of the line inputs are, but nothing about their relationship or the overall composition of the artwork. In future work, the regional perceptual logic layer will group line inputs into regions and containers to enable the system to learn and modify entire shapes (rather than individual lines). However, even without regional perceptual logic, the system was able to achieve complex (and artistically valuable) outputs in drawings 3-6 because the human starts defining themes and creating complex artistic patterns by drawing many lines per turn in rapid succession. The basic mimic functions of the Drawing Apprentice leveraged this complexity to achieve equally detailed output. The final product is shown in all black (as the artist saw it) in drawing 9.

To capitalize on the emergent nature of creativity in improvisation, our development efforts focus on building more sophisticated methods of perceiving, analyzing, and understanding drawn human input in such a way that it can be intelligently and creatively re-used by the system. This involves teaching the system how to recognize line groups (regional perceptual logic), how to define relationships between those line groups (global perceptual logic), and when it is appropriate to use them for generating artistic contributions (creative trajectory monitory).

In practice, the current prototype appears like a clumsy novice because it can achieve continuous improvisation, but it cannot detect patterns, make abstractions about the artwork, or understand any user intentionality. This limitation means that many of the system's contributions accidently disrupt things the user intentionally drew, such as a face or a nice curve. This creates tension for the artist and can serve as a creative catalyst or as a source of frustration if the disruptions are too severe or frequent. Skilled artistic collaborators are typically quite flexible and can integrate a wide variety of unexpected line contributions into their drawings with one key exception: completing the drawing.

When the artist was ready to complete the drawing by perfecting and refining each major line (drawings 7 & 8), the system kept blindly mimicking each line input, which effectively produced more work for the artist because each computer contribution was an unpolished line that required refining. This process eventually became frustrating because the artist wanted to stop but was never satisfied with the precision of the lines. Without global perceptual logic, the drawing as a whole cannot be evaluated to determines its level of completion.

With only the local and part of regional perceptual logic implemented, the Drawing Apprentice is able to maintain *continuous* collaboration with an expert artist, which is a milestone for the project. In addition to continuous collaboration, the final prototype will be successful if: (1) It provides similar benefits as a human collaborator (i.e. playful exploration, motivation, and creative engagement) (Carroll 2009); (2) Users find collaboration meaningful and valuable (Candy and Edmonds 2002); and (3) Implementing additional parts of the EMC increases creative engagement (Candy and Edmonds 2002).

Our research agenda includes a user study to evaluate the system. The study is a controlled experiment that compares collaborating with the Drawing Apprentice to human collaboration and a random control. Participants are asked to perform three collaborative drawing sessions on a tablet computer with an unknown 'player' as the computer collaborator (e.g. Apprentice, human, or random lines). After each drawing session, the participant will be interviewed and complete the Creativity Support Index to measure playful exploration, motivation, and creative engagement (Carroll et al. 2009).

## Conclusions

This paper described a cognitive model of enactive creativity that is useful for designing continuous improvisational collaboration in creative systems. We built an artistic computer colleague called the Drawing Apprentice to test our enactive model of creativity (EMC). The Drawing Apprentice embodies the principles of EMC using increasingly complex layers of perceptual logic to analyze and react to user input in real time improvisation. We hypothesized collaboration with computer colleagues based on the enactive model of creativity can enrich the creative process like human collaboration (i.e. increase playful exploration, motivation, creative engagement) in open-ended creative domain such as non-representational visual art. We presented the theory, design, prototype details, and early collaborative artwork generated with Drawing Apprentice, the co-creative drawing partner.

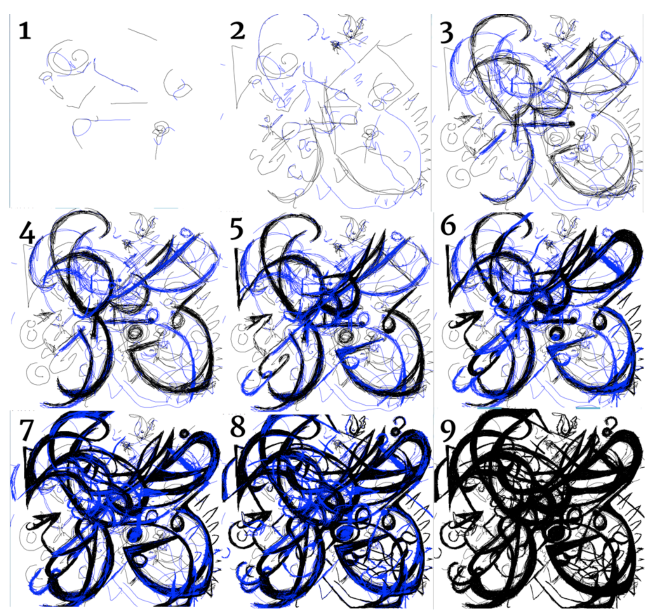

Figure 8: Time-lapse image of expert artist collaborating with the Drawing Apprentice (computer lines are blue).

## Acknowledgements

Thanks to Dr. Ellen Yi-Luen Do and researchers Cora Wilson and Monet Tomioka for helping build the Drawing Apprentice.

## References

Arnheim, R. (2001). *Art and visual perception*. Stockholms Universitet.

Adamson, R. E. (1952). Functional fixedness as related to problem solving: A repetition of three experiments. *Journal of experimental psychology*, *44*(4), 288.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and brain sciences*, *22*(04), 577-660.

Boden, M. A. (2003). *The creative mind: Myths and mechanisms*. Routledge.

Candy, L., & Edmonds, E. (2002, October). Modeling co-creativity in art and technology. In *Proceedings of the 4th conference on Creativity & cognition*, 134-141. ACM.

Carroll, E. A., Latulipe, C., Fung, R., & Terry, M. (2009, October). Creativity factor evaluation: towards a standardized survey metric for creativity support. In *Proceedings of the seventh ACM conference on Creativity and cognition* (pp. 127-136). ACM.

Clouzot, H., (1956). *Le mystère Picasso*. France: Filsonor.

Colton, S. (2008). Creativity Versus the Perception of Creativity in Computational Systems. In *AAAI Spring Symposium: Creative Intelligent Systems* (pp. 14-20).

Davis, N., Li, B., O'Neill, B., Riedl, M., & Nitsche, M. (2011, November). Distributed creative cognition in digital filmmaking. In *Proceedings of the 8th ACM Conference on Creativity and Cognition*, 207-216. ACM.

De Jaegher, Hanne. 2009. “Social Understanding through Direct Perception? Yes, by Interacting’ *Consciousness and Cognition* 18: 535-542.

Dourish, P. (2004). *Where the action is: the foundations of embodied interaction*. MIT press.

Gabora, L. (2010). Revenge of the “neurds”: Characterizing creative thought in terms of the structure and dynamics of memory. *Creativity Research Journal*,*22*(1), 1-13.

Gaspar, J. M., & McDonald, J. J. (2014). Suppression of Salient Objects Prevents Distraction in Visual Search. *The Journal of Neuroscience*, *34*(16), 5658-5666.

Gibson, J. J. 1986. *The ecological approach to visual perception*. Routledge.

Glenberg, A. M. (1997). What memory is for: Creating meaning in the service of action. *Behavioral and brain sciences*, *20*(01), 41-50.

Hoffman, G., & Weinberg, G. 2010, May. Gesture-based human-robot jazz improvisation. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 582-587. IEEE.

Kosslyn, S. M. 1996. Image and brain: The resolution of the imagery debate. The MIT Press.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine learning*, *1*(1), 11-46.

Li, B., Zook, A., Davis, N., & Riedl, M. O. (2012). Goal-Driven Conceptual Blending: A Computational Approach for Creativity. In *International Conference on Computational Creativity* (Vol. 10).

Lubart, T. 2005. How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies*, *63*(4), 365-369.

Magerko, B., Fiesler, C., Baumer, A., & Fuller, D. 2010, June. Bottoms up: improvisational micro-agents. In *Proceedings of the Intelligent Narrative Technologies III Workshop*, 8. ACM.

Mace, M. A., & Ward, T. 2002. Modeling the creative process: A grounded theory analysis of creativity in the domain of art making. *Creativity Research Journal*, *14*(2), 179-192.

Mendonça, D., & Wallace, W. A. 2004. Cognition in jazz improvisation: An exploratory study. In *26th Annual Meeting of the Cognitive Science Society, Chicago, IL*.

Nersessian, N. 2008. *Creating scientific concepts*. The MIT Press.

Newell, A. (1959). *The processes of creative thinking*. Rand Corporation.

Noë, A. (2004). *Action in perception*. MIT press.

Norman, D. A. (1999). Affordance, conventions, and design. *interactions*, *6*(3), 38-43.

Pressing, J. (1984). Cognitive processes in improvisation. *Advances in Psychology*, *19*, 345-363.

Stewart, J., O. Gapenne and E. Di Paolo. 2010. *Enaction: Toward a New paradigm for Cognitive Science*. Cambridge, Mass: The MIT Press.

Sawyer, R. K. 2012. Explaining creativity: The science of human innovation. Oxford University Press.

Schön, D. A. 1992. Designing as reflective conversation with the materials of a design situation. *Knowledge-Based Systems*, *5*(1), 3-14.

Shneiderman, B. (2007). Creativity support tools: accelerating discovery and innovation. *Communications of the ACM, 50*(12).

Varela, F.J, Thompson, E, and Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience.* Cambridge, Mass: The MIT Press.

Vygotsky, L. S. (1980). *Mind in society: The development of higher psychological processes*. Harvard University Press.