Toward a Cognitive Theory of Creativity Support

Nicholas Davis
School of Interactive Computing
Georgia Institute of Technology
ndavis35@gatech.edu

Holger Winnemöller, Mira Dontcheva
Adobe Systems, Inc.
{hwinnemo, mirad}@adobe.com

Ellen Yi-Luen Do
Industrial Design & Interactive Computing, Georgia Tech & CUTE Center @ IDMI, NUS
ellendo@acm.org

ABSTRACT
We present the beginnings of a Cognitive Theory of Creativity Support aimed specifically at understanding novices and their needs. Our theory identifies unique difficulties novices face and reasons that may keep them from engaging in creative endeavors, such as fear of failure, time commitment, and lack of skill. To test our theory, we use it to analyze existing creativity support tools from multiple domains. We also describe the design and initial implementation of a creativity support tool based on our theory. The creativity support tool, called StorySketch, is designed to empower storytellers without graphical skills to engage in visual storytelling.

Author Keywords
Creativity; Human Computer Interaction; Creativity Support Tools; Cognitive Science

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

General Terms
Human Factors; Design; Measurement.

INTRODUCTION
Creativity is a learnable skill that leverages common cognitive mechanisms [6]. Researchers in the field of Human Computer Interaction (HCI) have begun to develop tools that influence, support, and enhance creativity called Creativity Support Tools (CST) [19]. Most of these tools are aimed at creative professionals [19] as these are traditionally the individuals engaging in creative activities [16]. Some researchers argue that non-professionals do not engage in creative tasks because they are not capable of creative expression [1]. However, there is evidence that suggests that novices can be creative [16].

When it comes to supporting creativity, experts require tools that enable them to rapidly iterate through and document many different ideas early in the creative process [19]. Novices, on the other hand, require a low threshold of entry into the creative domain because they may not have the confidence, skill, or motivation to sustain engagement with the tool [19]. We present the Cognitive Theory of Creativity Support aimed specifically at understanding novices and their needs. Our approach helps identify potential difficulties for novices and reasons they may disengage with CST tools, such as fear of failure, time commitments, and lack of skill.

Evaluation metrics for professionally oriented CSTs typically measure creative performance with and without the tool [15]. However, CSTs that introduce completely novel experiences cannot be measured with this metric [15]. We argue that evaluating novice use of CSTs requires a new understanding of how CSTs affect cognitive processes.

In our analysis, we use cognitive theories to explain how CSTs can address the needs of novices. We use the cognitive science theories of embodiment, situated activity, and distributed cognition to build the concepts of *embodied creativity*, *situated creativity*, and *distributed creativity*. These three perspectives of creative cognition [6] form the foundation of the Cognitive Theory of Creativity Support and provide insight into how technology can be used to support the creativity of novices.

Embodiment describes how the human body—our sensory apparatus and how we interact with the environment—largely influences cognition [22]. The processes and mechanisms involved in creative cognition are therefore rooted in interaction. Embodied creativity outlines methods to make novice's creative ideas more concrete and interactive.

Situated Activity describes how routine actions are constraint-based improvisations with the environment guided by real time sensory feedback [1, 17]. Tool use in the creative process can detract from the situated nature of activity when novices are unaware of how to use tools to accomplish an action. Situated creativity describes a continuum of competency that shows how tools can become an extension of the body and allow individuals to focus on creative expression rather than consciously controlling tools.

Distributed Cognition proposes that humans can offload cognitive tasks onto other people, the environment, and tools to ease cognitive load [11]. Offloading creative tasks

C&C '13, June 17 - 20 2013, Sydney, NSW, Australia
Copyright 2013 ACM 978-1-4503-2150-1/13/06©$15.00.

that require significant skill onto CSTs can help novices produce richer creative output earlier in the creative process. Distributed creativity describes how automating technical skills to varying degrees can increase creative engagement, motivation, and reduce the barrier of entry for novices.

We begin by discussing different types of novices, their skillsets, audiences, and motivation for using CSTs. We examine how CSTs affect creative productivity and why expert-oriented CSTs present difficulties for novices. Next, we build the foundation of the Cognitive Theory of Creativity Support by examining Embodiment, Situated Activity, and Distributed Cognition in turn. Next, we use our theory to analyze CSTs from multiple domains. We identify potential difficulties for novices and use our theory to find solutions to these problems. Finally, we present the design process and initial implementation of a prototype in the domain of visual storytelling called StorySketch.

BACKGROUND

The first step we take in building our theory is defining terms that will help us articulate the unique needs of novices. We make a conceptual distinction between form and content. This form-content distinction helps us understand the skills involved in creative expression. Next, we make a categorical distinction between types of novices according to their skill creating form and content.

Delineating Form and Content

The form of an artifact is the physical shape or body of the thing being created [6]. The content is the meaning or expression the form signifies to the viewer [6]. For poetry and language arts, form corresponds to the individual words and their sounds, while content is the semantic meaning or narrative those words construct for the reader. For visual art, form corresponds to the shapes and the way those shapes are made including the lines or paint that comprise them, while the content is the image and the meaning derived from it, such as an emotion or idea. In many ways form and content are intimately connected. The style of a form may influence its interpretation. As creators become more skilled at their craft, they exercise greater control over the forms they use to convey their intended meaning. Artistic content becomes richer and more complex as an artist expands the range and nuance of forms they employ.

Anatomy of a Novice

In our theory, we distinguish between two types of novices—tool novices, and domain novices. Tool novices are familiar with the domain but they do not know how to use a CST to achieve their creative goals. Domain novices, on the other hand, are unfamiliar with the creative domain as well as the CSTs used to enhance creativity in that domain.

Tool novices are domain experts that have never used a CST in their domain before. These individuals know how to formulate and develop their ideas using non-digital tools, but cannot immediately map that knowledge to the CST due to unfamiliarity with the interface and functions. An expert oil painter, for example, is skilled at artistic techniques such as brushwork, color theory, and life drawing, but s/he may have no idea how to execute this expertise with the digital painting functions in Photoshop. We make a distinction between knowledge of how to execute an artistic idea and conceptual knowledge about the artistic idea [6]. The CST design goal for tool novices should increase the usability of the tool, including making the controls more intuitive and improving the amount of control that users have over artistic forms.

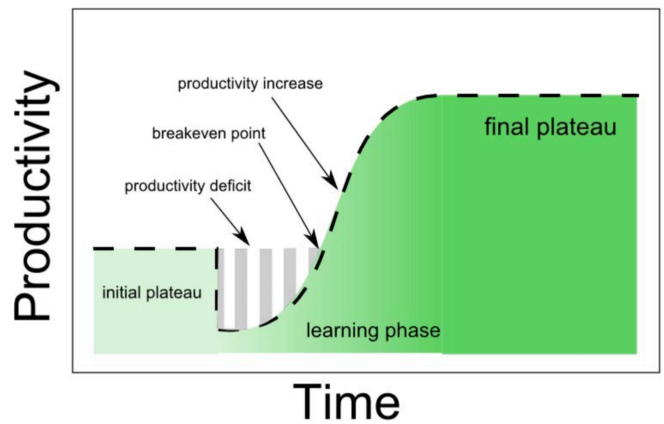

Figure 1: Time investment required to make CSTs useful.

Tool novices have the motivation and justification to invest the time required to learn how to use complex CSTs because they can substantially improve their creative output. Figure 1 represents the productivity deficit experienced during the learning process. While learning the tool, productivity is reduced because one is attending to the tool itself. At the breakeven point, however, the users productivity surpasses what was possible without the CST. Once learned, CSTs can allow individuals to rapidly explore, record, and organize many different creative ideas, which can significantly increase the quality and quantity of creative output [19].

The quality of creative products is of particular interest to tool novices because they are creative experts. Other experts in their field will evaluate their creative products according to their relative novelty and value [3,21]. For example, museum curators selectively chose artworks that will be displayed in galleries. Tool novices view CSTs as a mechanism to increase their standing in a field, which justifies the time learning a CST, even if it is a steep learning curve.

Domain novices are individuals that are unfamiliar with the creative domain as well as the CST tools associated with it. They are unaware of how to creatively express themselves in the target domain, i.e. they do not know how to create and manipulate artistic forms to develop meaning and realize artistic goals. Domain novices have no significant prior experience with the tools of the domain—digital or otherwise. In this circumstance, the CST enables a completely new creative experience that would be impossible without the support of the tool [15]. Domain novices use the CST for personal self-expression. The Cognitive Theory of Creativity Support can offer unique insights into this user group because they require a significantly different strategy to support their creativity. CSTs directed toward this group must have a "low barrier of entry" and help novices learn about the creative domain as they use the tool [19].

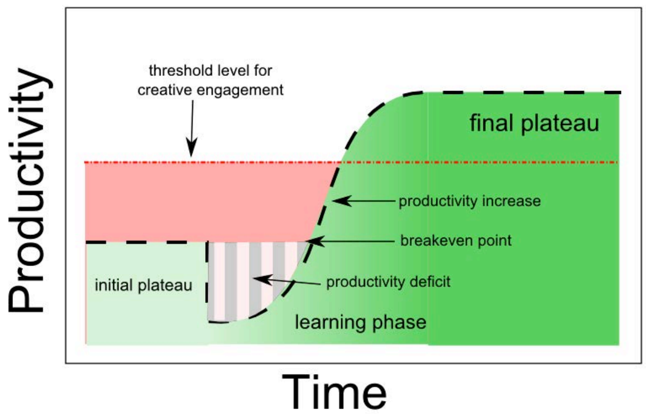

Figure 2: Time investment required for creative engagement with CST.

Domain novices may not be as willing to invest a large amount of time learning the tools or the domain because they are casual creators. When a domain novice engages a CST designed for an expert, they may become discouraged because they do not see the desired results early on. Figure 2 shows how creative engagement with a CST only occurs after a significant time investment learning the controls. However, even once domain novices learn the tools, they may not be able to effectively create artistic content because they don't know how to combine artistic forms to realize their goals. Because of their unwillingness to invest time in learning how to use CSTs, domain novices or "casual creators" are underrepresented within the CST community [15].

Distinguishing between tool novices and domain novices allows us to focus on the unique needs of domain novices, which include teaching both the tool and the domain. To provide insight into how technology can enhance and teach creativity to novices, we now turn to cognitive science.

COGNITIVE THEORIES

Cognitive science theories describe the way we think, perceive, build concepts, and learn [22]. Here, we will examine three modern cognitive theories to help us understand how technology affects the creative process.

Embodied Cognition

The embodiment theory argues that the human body structures cognition in many important ways [22]. Cognition is not the disembodied and mechanical manipulation of symbolic representations that traditional cognitive science theories posit. Instead, human cognition is rooted in perception, action, and interaction with the environment [23]. We propose a subcategory of embodied cognition referred to as *embodied creativity*, which claims that the cognitive mechanisms involved in creativity, such as idea generation, conceptual combination, exploration, visualization, mental imagery, and problem solving/finding are dependent upon and facilitated by interaction with the environment. In other words, creativity is fundamentally rooted in interaction with the world.

The importance of sketching in many creative disciplines supports our claim that creative cognition is embodied. Schon famously describes sketching as a conversation with one's creative materials [20]. It is a low cost and flexible exercise that helps artists, designers, engineers, and many other creative domains explore and develop ideas [22]. Given that sketches can vary considerably based on the domain, it is useful to develop a general definition of a sketch as a visualization that exists on an external medium (not in the head), which represents an idea or concept. For example, architects construct models to help them conceptualize the spatial properties of a building design. Models are physically very different from a hand-drawn sketch, but they still help offload the visualization of an idea, and therefore serve a similar cognitive function in the creative process as a hand-drawn sketch. Collages, storyboards, and mood boards are all different techniques of developing an external representation of an idea. Although they correspond to different creative domains, the general tendency to use sketches to externalize an idea is pervasive.

The visual relationships individuals detect from inspecting a sketch vary based on the domain as well, but there is a common iterative cycle of perceiving and manipulating the sketch. The discoveries made through this reflective and interactive process reciprocally shape and are shaped by one's creative goals [20]. That is, the idea of what one would like to create changes and evolves as one begins to see the idea materialize and interact with it [20]. The design goal that follows from this analysis is that CSTs should reduce the time it takes to make ideas concrete and interactive. This is particularly relevant to domain novices because they lack the specialized execution knowledge that is commonly required for sketching and rapidly prototyping ideas.

Physically changing the representation of an idea by tweaking its appearance is also possible once it exists as an external representation. Line sketches are particularly amenable since they are low-cost and easy to produce, which is one reason they are so useful in the creative process. Another design goal to support embodied creativity with CSTs is lowering the cost of making physical changes and increasing the interactivity of the system. Embodied cognition tells us that we think through action, and our CSTs should therefore be designed to facilitate this interactive creative thinking.

One dimension of interaction is altering the way in which the representation is viewed to explore new interpretations. For example, physically changing the viewing angle, distance, and scale of the artifact can reveal new relationships. In visual art, for example, painters often take a step back from their painting to gain a new perspective. The different scale allows a more global perception of the artwork, which reveals new relationships and offers new fodder for analysis not visible up close. In other domains, such as writing, there are analogous procedures, such as creating an outline or storyboard to view a summary of events at different scales. The level of visual abstraction as well as the physical distance used to represent information can change how that information is processed. Embodied creativity shows us that creative users need to be able to control these features fluidly.

Situated Cognition

Situated Activity (SA) argues that the environment is "a rich source of constraints and opportunities" that provide context to actions [1]. Without the environment, actions are unarticulated, generic, and abstract notions in the mind. Activity is situated because it is an improvised interaction with the environment guided by real time feedback.

This philosophy of action has its roots in Heidegger's phenomenological insight that objects can be *ready-to-hand* or *present-at-hand* [12]. In daily life, "we usually encounter things as resources for immediate action in the service of achieving our goals" [1]. For example when driving a car, the steering wheel in some sense ceases to exist. The steering wheel becomes an extension of the arms and is ready-to-hand. One does not need to consciously plan steering; activity occurs through small automatic adjustments based on feedback from perceiving the road. However, when a problem arises the steering wheel becomes present-at-hand and emerges from the undifferentiated background of objects to be consciously controlled or inspected based on its properties [1]. Expert drivers are able to automatically use small cues, such as the subtle behavior of other cars on the road, to adjust their current heading.

One of the goals of developing a Cognitive Theory of Creativity Support is to be able to have a systematic design strategy that helps us build CSTs that novices will be able to use in a way that is ready-to-hand. To achieve this, we require a more nuanced description of what Heidegger meant when he described the phenomenological transition from present-at-hand to ready-to-hand. Vygotsky's Activity Theory will be used in this analysis because he developed a detailed vocabulary to refer to actions in the process of becoming an expert [13]. He defined three hierarchical components of an activity shown in Figure 3.

Figure 3: Vygotsky's activity triangle with Heidegger's tool descriptors

At the highest and most abstract level is the *activity*, which is the motivation, intention, and reason for engaging in the activity [13]. The middle level consists of conscious *actions* that all have their own goals that relate to the object [13]. When one consciously executes an action, the tool is present-at-hand. Finally, the lowest level of the hierarchy consists of *operations*, which are subconscious routines that implement actions according to the conditions of the environment [13]. Routine operations employ tools in a ready-to-hand manner guided by feedback from the environment.

Figure 4 shows the process of a tool novice learning how to map knowledge to a new medium. Continuing with our previous example, we assume that new drivers know the rules of the road, but not how to execute that information. As the individual becomes expert, more actions are automated as operations, which is represented by the increased size of the operation part of the triangle. The rightmost triangle represents an expert that experiences steering as ready-to-hand.

Vygotsky called the gradual learning process that shifts a tool from present-to-hand to ready-to-hand the *broadening scope of action* [13]. It describes a process by which one can gradually acquire expertise by automatizing the execution of actions and thereby shift cognitive resources towards higher-level constructs, such as the overall artistic intentions and style of an artwork. As novices experiment with tools, they begin to create a cognitive model about how their actions relate to the objectives of a domain. Conscious actions become routinized as operations as their internal model is validated. The tool then starts to become ready-to-hand [13]. Instead of concentrating on using the tool, one is concentrating on the activity itself [13]. Here, executing operations becomes a situated and improvised activity that is automatically adjusted based on the current conditions of the environment. When the tool recedes into the background and the individual is fully absorbed in the activity, they experience what Csikszenmihalyi refers to as

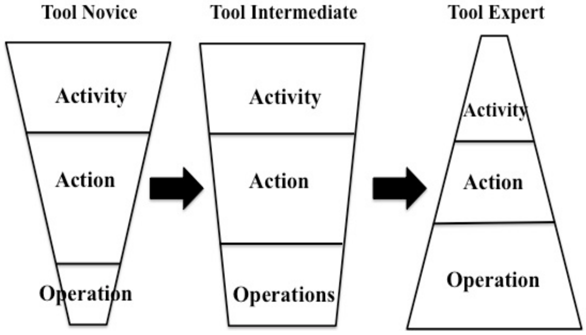

Figure 4: A progression from tool novice to tool expert depicted by Vygotsky's activity triangle.

flow, which is an optimal experience that correlates to more creative outcomes [4].

The learning process of tool novices can be modeled using the broadening scope of action because these users know what they are trying to achieve—they just have to learn how to master the tools. After they understand the mapping between their artistic knowledge and the functions of CST interface, their actions will gradually become automatic operations.

The learning process for domain novices, however, does not follow the standard broadening scope of action because domain novices lack two distinct types of knowledge—execution knowledge about the tools and content knowledge about the artistic forms those tools help create. Novices do not understand how creative actions relate to their artistic intentions, nor how to execute creative actions with a given toolset. These problems can retard progress from the novice to intermediate stage as shown in Figure 4. To bridge this gap, domain novices require CSTs that systematically offload some of the technical and conceptual skill required to create artistic forms.

Distributed Cognition

Distributed cognition (DCog) describes how humans can ease their cognitive load by offloading tasks onto the environment, tools, and other individuals [11]. Sketching allows individuals to offload cognitive processes, such as mental visualization, and frees cognitive resources for other creative reasoning tasks [20]. It facilitates spatial reasoning, reflection, exploration, analysis, and evaluation in the creative process. We have created the new term distributed creativity to describe the critical enabling function that CSTs like sketching play in offloading cognition in creativity. Rapidly externalizing ideas stimulates and enables individuals to engage different kinds of cognitive processes that support creative thinking.

Unfortunately, novices may lack the technical skill to produce sketches that are sensible and provide insight. The technical skills required to externalize ideas in a new domain can provide a significant barrier to distributed creativity. The individual may be able to vaguely imagine the content they wish to illustrate, but creating the proper forms to externalize that idea can represent a bottleneck that prevents the individual from exploring, developing, and refining their ideas. We propose that there are two basic skills required for distributed creativity:

- **Technical Skill:** Procedural knowledge about how to execute a task; the individual has to have a certain physical dexterity to render visual forms.
- **Conceptual Skill:** Semantic knowledge about what task to execute; the individual has to have a clear idea about how to manipulate and arrange visual forms to express meaning.

Tool novices may have the non-digital technical drawing skills, but they do not know how to map those skills to a drawing CST. Domain novices do not possess technical or conceptual skills, and we posit that CSTs targeted towards these users should be designed to offload both of these skills to varying degrees at different points in the learning process.

Scaling Distributed Creativity

Figure 5 shows a sliding scale of distributed creativity that describes how the technical and conceptual skills required to rapidly create and explore artistic forms can be offloaded onto CSTs to different degrees. The y-axis represents the creative output of an individual. The x-axis represents the degree of automation, or how much technical and conceptual skills are offloaded onto CSTs. When a domain novice approaches a tool with no automation, they are hardly capable of any creative output. As the degree of automation rises, so too does the creative output. When automation is 100%, the computer would be generating procedural art without any human input whatsoever. Although the user would see a pleasant image, there is little to no creative engagement in watching the computer generate art using AI algorithms. As automation decreases, the user has more control over the creative output, which increases creative engagement and motivates the user.

The concept of distributed creativity can help CST designers understand and leverage the cognitive benefit of sketching. We have identified two independent skillsets,

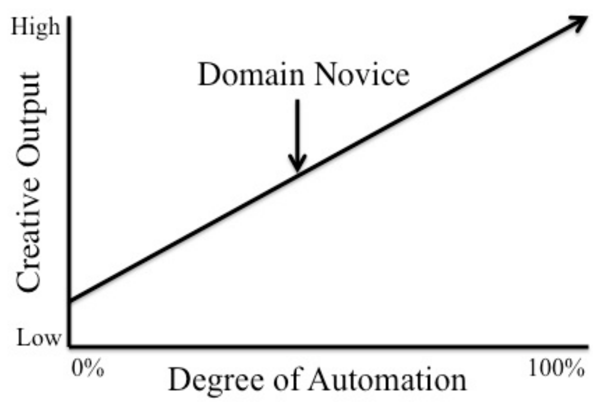

Figure 5: Scaling distributed creativity to achieve higher creative output for domain novices.

technical and conceptual skills, which can be offloaded to varying degrees. In this next section we will explore the implications of using CSTs to scale the degree to which these skills are offloaded. This discussion informs the existing CST design goal to enable low cost and flexible sketch-like content creation that allows one to externalize ideas [24].

First let us consider modulating the degree of automation for conceptual skills in the domain of visual art. Conceptual skills in this domain include knowing how to manipulate and arrange shapes and colors to convey an intended meaning or effect. Automating this process could include providing a set of *artistic primitives* that serve as basic building blocks that frame the conceptual structure of an artwork [17]. These artistic primitives could be combined and arranged in different ways to explore artistic ideas. A coloring book is an example where the conceptual development and layout of an artwork is provided while the user is responsible for the technical execution of coloring each region. At high degrees of conceptual automation, the artistic primitives would be large and could include entire characters, backgrounds, objects, etc. As the degree of conceptual automation decreases, the user would have more control over how those forms appear and are arranged, i.e. changing the hair color of characters or the amount of trees in the background. Examples of even more fine-grained artistic primitives would be the shading of a particular region of a character.

Next, let us consider modulating the degree of automation for technical skills in the domain of visual art. To automate technical artistic skill, a system can abstract away the physical dexterity required to execute an artistic idea, such as the process of character creation or the act of shading. At high degrees of automation, the commands remove all artistic skills, such as drag-and-drop or point-and-click actions where entire characters can be inserted through menu selection. Reducing technical automation would change the method of creating artistic forms to more natural mappings, such as the embodied actions involved in paper and pen drawing. For example, instead of dragging and dropping pre-fabricated characters, users could sketch a rough articulated skeleton to determine the position and pose of characters or the shape of props (see the section describing the StorySketch prototype for additional details on this input method). In this way, the user can begin to understand what it feels like to create the rough physical proportions of a character without being responsible for all of the detail.

Offloading technical and conceptual skills onto CSTs provides a learning scaffold that may help users gradually learn to be more creative. Scaling distributed creativity enables novices to develop and control creative objectives from a high-level, which allows them to explore and become familiar with the domain. Higher degrees of computer automation ensure that domain novices will enjoy success and positive feedback early in the learning process without a significant time investment. As users become more comfortable with the conceptual and technical knowledge of a creative domain, they can begin to reduce the degree of automation. This allows them to control finer grained details and create more nuanced forms to develop complex compositions.

Scaling distributed creativity allows novices to rapidly explore general creative decisions at a high level, which Vygotsky refers to as the activity level. Instead of struggling with the technical details of a new domain, novices that are introduced to cognitive-based CSTs (i.e. CSTs that align with the cognitive principles we describe here) would be able to explore a creative domain right away. In these early learning stages, the computer would generates all the requisite forms and a lot of the conceptual structure of an art piece while the user evaluates these forms through interactive creative exploration using abstracted commands. Once the user begins to understand all the variables of the creative domain and how they relate to each other, s/he develops a mental model about the conceptual rules of the domain. For example, in visual art novices may have to familiarize themselves with color theory, the use of shading to demonstrate volume and depth, and composition. Equipped with a rough mental model about these domain variables, the user can then choose to reduce the automation and begin exploring with different forms of technical execution and conceptual structure.

Scaling distributed creativity allows domain novices to accomplish expert-level creative output early in the creative process by automating some of expert actions and operations. Figure 6 shows an expert activity triangle with a bar representing the degree of control for which the user is responsible. When the user begins interacting with a highly automated CST, the bar is high, which corresponds to a high degree of automation. This means domain novices are able to control the artistic forms from a higher level, i.e. through textual parameters or pre-set configurations. As

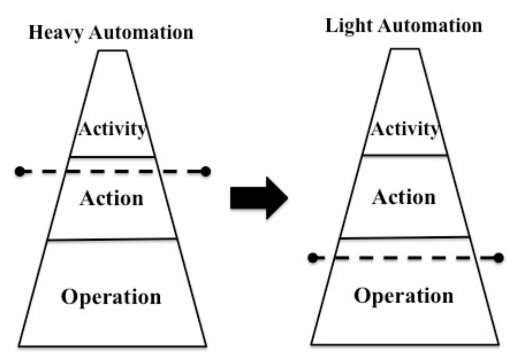

Figure 6: Scaling distributed creativity depicted with Vygotsky's activity triangle. The area above the bar shows what component of activity the user is responsible for.

automation decreases, more actions and operations are revealed to the user. These new actions enable more nuanced control over artistic forms, which results in richer creative expression.

CREATIVITY SUPPORT TOOL ANALYSIS

Thus far, we've laid the groundwork for a Cognitive Theory of Creativity Support. The foundation of this theory describes how embodied, situated, and distributed creativity provide us with design insights that can help novices. Together, these theories show that creativity is an interactive process that can be partially distributed upon technology. The way in which creativity is distributed can benefit novices by increasing the ways in which they can create, interact with, and evaluate their ideas. To further elucidate this theory and demonstrate its descriptive utility, the next section will employ its principles to analyze CSTs from several artistic domains.

Computer Programming

Computer programming is a creative activity that has a high barrier of entry for beginners. The complex syntax of programming languages requires a tremendous amount of technical knowledge. Conceptually, programming requires logic and an understanding of how to formally describe problems algorithmically. In this section we analyze an innovative programming language called Processing designed to support and teach creative programming skills.

Processing (processing.org) is a programming language and interactive development environment (IDE) designed for visual artists that are not necessarily familiar with computational concepts. It supports embodied creativity because it reduces the amount of time required to visualize ideas using code. The language is meant to help designers 'sketch' out ideas using code, which is evident by the ease with which a canvas and visual elements are created. For example, Processing only requires 1 line of code to draw a line to a canvas whereas Java requires approximately 50.

The exhaustive list of interactive examples on the Processing website demonstrate the power of programming, which helps novices build a conceptual framework about what can potentially be created with Processing. Learning about the affordances of a tool supports situated creativity; it helps novices understand what class of problems the tool is appropriate for.

The type of example-based learning that Processing promotes supports distributed creativity as well. Novices can copy and paste entire examples and tweak them to learn about all the working parts and parameters that generate the procedural artwork. Users can spend their time exploring different ways to create neat looking visualizations with code rather than figuring out the cumbersome syntax to create the visual forms in the first place.

The major pitfall of Processing is its inability to scale distributed creativity. There is a significant jump in skill required between *modifying* and *adapting* examples. Modification occurs in small increments, but heavy adaptations or combining multiple examples can be extremely difficult without thoroughly understanding the source code. Comments in the source code can help explain what each section does, but the terminology used will most likely be difficult for a beginner to parse. Acknowledging that example based programming should be a continuum and bridging the gap between modification and adaptation will support domain novices approaching Processing as a CST. For example, domain novices may know that examples x, y, and z have the functionality they require, but they may not know how to write the code that combines these functions. One potential solution that may be helpful is if domain novices could select from a list of functions and have a CST system adapt these examples into a coherent fully coded infrastructure. This would bridge the gap between modification and adaptation, and create a more continuous learning scaffold.

Music

Apple's Garage Band is a program for both the desktop and tablet that allows users to create, record, and mix music. It provides an excellent example of a cognitive creativity support tool because it scales distributed cognition and helps novices quickly create aesthetically pleasing music. A beginner can sit down with the program and click (or tap) through a few chords to make a pleasant sounding riff. While users play and experiment, they can record what they are creating and play it back, edit it, or layer other instruments on top of it. The initial time investment for creating a harmonious song is significantly reduced, which supports the interactivity requirement of embodied creativity.

This CST has a robust method for scaling distributed creativity. Novices can make music on the device by playing 'smart instruments' that reduce the technical skill required to play pleasant sounding chords. The chords displayed on the screen are also curated using music theory to ensure that they sound good together in any combination, which is an example of offloading expert conceptual knowledge.

As novices progress, real musical instruments can be plugged into the computer, such as electronic keyboards and guitars. The program offers guided lessons for these instruments. The user learns by playing along with existing songs. This offloads conceptual skills because the structure of the song already exists, and it offloads technical skills because it steps the user through the music. The CST scales distributed creativity by providing lessons with real instruments as a midpoint between heavy automation (smart instruments) and no automation (freeform jam with real instruments).

Storytelling

In this section, we describe a prototype we are currently developing, StorySketch, which enables casual creators to quickly experiment with visual narrative ideas and work iteratively and interactively as a means of self-expression. StorySketch allows individuals to creative visual narratives using flexible and low cost sketches.

StorySketch is designed as a cognitive creativity support tool in the sense that it is a prototype based on the Cognitive Theory of Creativity Support. The motivation and design rationale for the tool was to create a fun and engaging way for individuals without artistic skills to creatively express themselves through storytelling. Given that storytelling is an important mode of creative expression, we wanted to empower individuals to tell their stories visually.

The central design consideration was offloading the technical skill required to create aesthetically pleasing visual forms, such as a characters and props. We wanted users to spend their time developing their stories rather than having to figure out how to draw all the characters and props for their story. StorySketch should allow individuals to quickly explore different visual storytelling strategies and evaluate which style works best for their particular story. The key design challenge was how to create an input method that allows rapid and nuanced control of character and prop creation. The input method should be relatively easy given that our target audience is domain novices that do not necessarily have technical sketching skills.

One solution for enabling users to rapidly explore visual forms is dragging and dropping static forms. This solution uses the artistic primitive idea discussed in the distributed creativity section. It reduces the technical skill necessary for visually representing ideas, but it also limits the complexity of information that can be communicated to the computer. The user can only determine the position of story elements. Making characters dynamic and malleable could address this issue. For example, the limbs could be dragged into the desired position. However, the character creation process in this solution is now at least two steps and perhaps more given that each limb has to be adjusted independently. Additionally, this approach reduces the richness of the creative experience because users may feel disconnected from the characters since they did the input method was completely disembodied.

Our solution to this problem was to use a naturalistic yet simplified version of drawing story assets. For example, when creating a character, users draw a stick figure as the input method. The strokes and general form of the input is analyzed by the system to create an articulated character rendered by a physics engine as shown in figure 7. A character model is then applied to the articulated character as shown in Figure 8. This solution leverages a very basic degree of technical skill to generate a character that is sufficiently complex in terms of pose and position.

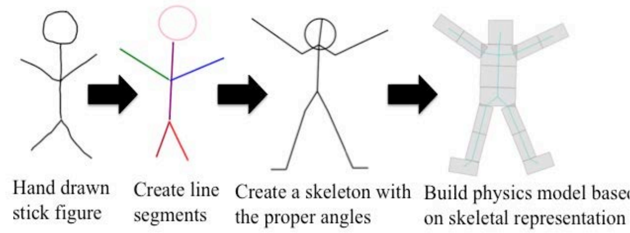

Figure 7: Sketch based character creation

The form of the skinned character still corresponds to the user's sketched stick figure, but the exact details and technically difficult part is handled by the system, such as proportions, shading, and coloring.

This approach is amenable to future extension by gradually increasing the amount of user control during the creation phase. The user could sketch customizations, such as facial expressions, cloth colors and folds, and shading. This type of interaction provides a continuum of distributed creativity to scaffold the user's skills and creatively engage the user. The user should be free to define the amount of work they would like to do and the amount of work offloaded onto the computer. It is conceivable that an individual may not be interested in customizing characters, but would rather choose to spend their time exploring different storylines.

The sketch-to-character interaction supports the creation and exploration of content, but problems arise with respect to evaluation. Evaluation comes in two parts—individual and social. Does the story match the user's intended goal? Does the story effectively communicate the intention of the author?

For StorySketch to help create the user's imagined mental content, the character models the system uses to assist in character creation must directly relate to the user. Initially, we used a stock cartoon character in formative user testing. However, a relatable and diverse set of characters quickly became the number one request of users. How can users tell a meaningful and personal story about a specific set of people with generic characters?

From a design standpoint, this problem represented an

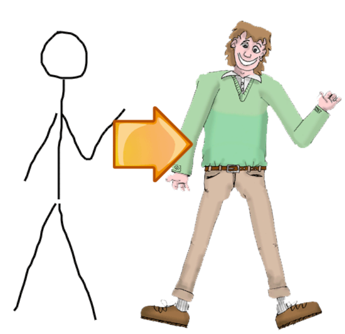

Figure 8: Fully skinned character

authoring bottleneck—who should bare the responsibility for creating user-specific characters? The user can't manually create these characters because they are presumed to not have enough artistic skill. The designers can't possibly create a wide enough set of characters to remain relevant to all users. One solution is to use the avatar model used in video games such as the "Mii" characters that users create on the Nintendo Wii. However, manual character personalization requires a significant time investment, which can further delay creating narrative content. This may be a viable option to gain more control over forms, but there should initially be a way to control artistic forms at the most abstract activity level. Creativity support tools should strive to reduce the time required for externalizing ideas. Our planned solution allows users to import photos of their friends to create characters. The system identifies the target person in the photo, applies a cartoon filter, and then creates an interactive character. This character can be customized using a similar sketch based interaction as the current prototype.

This discussion of StorySketch demonstrates how our theory was useful in the design process of a cognitive creativity support tool targeted towards domain novices. Our theory can provide a descriptive framework for analyzing existing technologies as well as offer prescriptive guidelines for supporting creativity in novices.

DISCUSSION

Dynamically adjusting the amount of support and automation provided to domain novices as they engage with a tool and increase their creative skills can help them achieve the optimal creative experience, or flow [4]. Csikszentmihalyi found that individuals enter 'the zone' and are most creative when they are confident in their skills yet challenged by their task at hand. The proposed theory outlines how CSTs can provide a sliding scaffold that distributes some of the creative skills novices may lack, which include technical skills representing artistic forms and a conceptual understanding of how those forms should be arranged and manipulated. By adjusting the skill required, cognitive creativity support tools can help users achieve flow and increase creative output.

The Cognitive Theory of Creativity Support is meant to be domain independent; it doesn't outline the exact method or mechanism by which one should offload technical or conceptual skills because skills vary between domains. In the case of Garage Band, the tool offloaded the conceptual knowledge of music theory and the technical skill required to operate instruments. This support system was stepwise; a continuous scale (as we showed in Figure 4) might be overly optimistic or unrealistic because different technologies, interfaces, and input methodologies may be required to offload different types of skills to varying degrees. For example Garage Band changed interface when moving from 'Smart Instruments' to music lessons on real instruments. The technology behind each of these tools is significantly different as well. The StorySketch prototype proposed a more continuous example showing how sketches could become gradually more detailed by allowing more nuanced input commands, such as sketching facial features.

There is an implicit assumption in situated and distributed creativity that domain novices would like to gradually gain expertise and increase their skill in a domain. However, it may be the case that domain novices would simply like to generate creative content without increasing their technical or conceptual understanding of a domain. Increasing the degree of control over artistic forms may overcomplicate the tool for these users. Not all users want to learn to be more creative—they may just want to create cute, funny, or pleasing art projects with ease. For these novices, it may be best to find a 'sweet spot' on the scale of distributed creativity that offloads an optimal amount of skill.

Our conception of 'creative output' in the paper is necessarily vague. Creativity is usually defined as being novel, valuable, and sometimes surprising [21]. However, that definition is implicitly social—it describes a socially negotiated sense of value and novelty as compared to a field that contains other creative products. Domain novices don't necessarily mind if their creative products aren't hailed as 'genius.' Creative output refers to a more personally defined sense of little-c creativity that would be valuable to the user and his or her social circle.

Future work will expand upon the Cognitive Theory of Creativity Support and subject it to additional evaluation to determine its validity and utility. One evaluation approach is to use it as a design guide and perform user studies on the cognitive creativity support tools developed based on its principles. When evaluating these tools, we suggest that researchers carefully consider the qualitative and subjective user experience. Researchers may find the Creativity Support Index (CSI) valuable to measure the subjective quality of user's experience as they engage with CST [2]. Subjective experience is important because it impacts motivation and the users perceived level of creativity, which can increase the time that they spend working on creative endeavors. The CSI can help evaluate the effectiveness of some of the tradeoffs identified by our theory, such as the ability to quickly generate content versus the amount of control over that content or the degree to which technical or conceptual skills are automated.

CONCLUSION

We have presented a Cognitive Theory of Creativity Support that uses cognitive theories to identify the unique needs of novices. Embodied creativity shows how creativity is inherently interactive—technologies that help novices express and interact with creative ideas can increase motivation and creative engagement. Situated creativity describes how tools can both help and hinder creativity—it identifies a continuum between consciously focusing on a tool and automatically using that tool as an extension of the body based on real time feedback of the creative task. Distributed creativity describes how CSTs can offload some of the conceptual and technical skills required to use tools and artistic forms for creative expression. Distributed creativity can facilitate situated and embodied creativity by offloading technical and conceptual skill required to create and interact with ideas throughout the creative process. Together, these three concepts comprise the beginnings of a Cognitive Theory of Creativity Support. We used this theory to analyze existing CSTs to find their strengths and weaknesses with respect to novices. The theory was also used to guide the design and implementation of a CST designed to support visual storytelling called StorySketch. Future work can build on this emerging theory to discover additional techniques to support and teach novices how to be creative.

ACKNOWLEDGMENTS

This research is supported in part by Adobe Systems, Inc., the US National Science Foundation under Grant No 1117665 (SHB) and the Singapore National Research Foundation under its International Research Center Keio-NUS CUTE Center @ Singapore Funding Initiative and administered by the IDM Program Office. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of the sponsors.

REFERENCES

1. Beer, R. Dynamical systems and embedded cognition. In K. Frankish and W. Ramsey, eds., *The Cambridge Handbook of Artificial Intelligence* (856–873). Cambridge Univ Press, Cambridge, 2010.
2. Carroll, E.A., Latulipe, C., Fung, R., and Terry, M. Creativity factor evaluation: towards a standardized survey metric for creativity support. *Proceedings of ACM Creativity & Cognition*, (2009), 127–136.
3. Csikszentmihalyi, M. Implications of a Systems Perspective for the Study of Creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (p. 313). Cambridge Univ Press, 1999.
4. Csikszentmihalyi, M. *Creativity: Flow and the psychology of discovery and invention*. Harper perennial, 1997.
5. Davis, N, Li, B., O'Neill, B., Riedl, M., & Nitsche, M. Distributed creative cognition in digital filmmaking. *Proceedings of the 8th ACM Creativity and Cognition* (2011), 207–216.
6. Davis, N., Zook, A. Riedl, M, Kirschner, F., Nitsche, M. Techniques for Evaluating Novice-Oriented Creativity Support Tools. *Proceedings of CHI 2013 Workshop on Evaluation Methods for Creativity Support Environments* (2013).
7. De Saussure, F. *Course in general linguistics*. Columbia University Press, 2011.
8. Finke, R. A., Ward, T. B., & Smith, S. M. *Creative cognition: Theory, research, and applications*. MIT press Cambridge, MA, 1992.
9. Gibson, J. J. *The ecological approach to visual perception*. Hillsdale, Lawrence Earlbaum, 1979.
10. Hennessey, B.A. The consensual assessment technique: An examination of the relationship between ratings of product and process creativity. *Creativity Research Journal 7*, 2 (1994), 193–208.
11. Hutchins, E. (1995). *Cognition in the Wild*. MIT press Cambridge, MA.
12. Ihde, D. Heidegger's Philosophy of Technology. In R.C. Scharff and V. Dusek, eds., *Philosophy of technology: the technological condition: An Anthology*. Blackwell Publishing, Malden, 2003.
13. Kuutti, K. Activity theory as a potential framework for human-computer interaction research. *Context and consciousness: Activity theory and human-computer interaction*, 1996, 17–44.
14. Larkin, J., & Simon, H. Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, *11*, 1 (1987), 65-100.
15. Nakakoji, K. Meanings of Tools, Support, and Uses for Creative Design Processes. *Proceedings of the International Design Research Symposium'06*, 2006, 156–165.
16. Reilly, R. C.. Is expertise a necessary precondition for creativity?: A case of four novice learning group facilitators. *Thinking Skills and Creativity*, *3*,1, 2008, 59–76.
17. Resnick, M., Myers, B., Nakakoji, K., Shneiderman, B., Pausch, R., Selker, T., & Eisenberg, M. (2005). Design principles for tools to support creative thinking. *Creativity support tools: Report from a US National Science Foundation sponsored workshop*.
18. Suchman, L. *Human-Machine Reconfigurations: Plans and Situated Actions (Learning in Doing: Social, Cognitive and Computational Perspectives)*. Cambridge University Press, 2006.
19. Schneiderman, B. Accelerating Discovery and Innovation. *Communications of the ACM*, *50*, 12, 2007.
20. Schön, D.A. Designing as reflective conversation with the materials of a design situation. *Research and Engineering Design 3*, 3 (1992).
21. Sternberg, R. J. (2006). The nature of creativity. *Creativity Research Journal*, *18*(1), 87–98.
22. Thagard, P. *Mind: Introduction to cognitive science*. The MIT Press, 2005.
23. Varela, F. J., Thompson, E., & Rosch, E. *The embodied mind: Cognitive science and human experience*. MIT press, 1999.
24. Voigt, M., Niehaves, B., & Becker, J. Towards a Unified Design Theory for Creativity Support Systems. Proceedings of DESRIST'12 (2012), 152–17.