

An Artistic Dialogue with the Artificial

David Norton, Derrall Heath, Dan Ventura

MIND Lab
Computer Science Department
Brigham Young University
Provo, Utah
USA

dnorton@byu.edu, dheath@byu.edu, ventura@cs.byu.edu

ABSTRACT

In conjunction with Brigham Young University's Visual Arts program, we conducted a study centered around a system designed to be an artificial artist, in order to synthesize the ideas of visual artists and computer scientists. Participants from both disciplines designed activities that imposed the limitations of the artificial system on their fellow participants. These activities sparked discussion and insight into the nature of the creative process and how it can be better emulated in artificial systems. We present our system and several of the activities designed around it and discuss the synergistic results.

Author Keywords

creative process, art, collaboration, evolutionary system

ACM Classification Keywords

I.2.m [Artificial Intelligence]: Miscellaneous; K.3.1 [Computers and Education]: Computer Uses in Education — collaborative learning

General Terms

Design, Documentation, Experimentation, Human Factors

INTRODUCTION

In order to instigate a dialogue, a group of college art students was presented with a series of anonymous images and asked to comment on the creativity and value of the images. The images ranged from Van Gogh's *Starry Night* to Malevich's *Red Square* to a fourteen year old student's meticulous reproduction of *Starry Night*. They included a painting titled *Purple Haze* by Rumba the Wonder Horse, a digital image from the Mandelbrot set, and a phase contrast micrograph of *Trichodina pediculus*. They also included three images each produced by one of the digital artists: AARON [?], The Painting Fool [?], and DARCI [?]. The students knew the work of Van Gogh and Malevich. They could immediately

spot the fraud of the *Starry Night* reproduction. They marveled at the mathematical beauty of the Mandelbrot set and phase contrast micrograph. None of them knew the works of Rumba the Wonder Horse, or the three digital artists—while intrigued, they did not know how to appreciate these works. From these students' perspective, they could not properly attribute creativity without understanding the process behind the creation.

This exercise inaugurated a four month study designed to create synergism between the (apparently disparate) disciplines. With the collaboration of students and faculty from Brigham Young University's Visual Arts department, we designed the study to explore the process of creation in the visual arts. Our interest in the collaboration was to learn how to better emulate the human process of creation in our digital artist, DARCI. For the participating art students, it was an opportunity to explore their own personal approaches to creation and why they do what they do when they create.

The centerpiece of the study was DARCI (Digital ARTist Communicating Intention) a computer system designed to create original images through processes perceived as creative. We facilitate the perception of creativity by giving the system a name and referring to DARCI as a she. *She* is central to our research in the budding field of computational creativity. Computational creativity is a branch of artificial intelligence focusing on emulating human creativity in computer systems. One of the goals of this field is to automate tasks that one would say require creativity to perform such as mathematical deduction, the creation of art and music, and storytelling. Another goal of the field is to better understand creativity so that we may identify, foster, and augment the creative process in society.

DARCI presently does not produce wholly original images; rather, using a variety of image processing filters, she renders provided images so that they will reflect an accompanying description, currently limited to adjectives (see Figure ?? for example). In accordance with Colton's arguments concerning *appreciation* [?], this process requires that DARCI be able to evaluate the degree to which an image matches a list of adjectives. She is able to make this assessment by learning to associate image features with individual adjectives via supervised machine learning. In other words, the data used to train DARCI is created by teachers—humans.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

C&C'11, November 36, 2011, Atlanta, Georgia, USA.

Copyright 2011 ACM 978-1-4503-0820-5/11/11...\$10.00.




Figure 1. Example of DARCI in action.

The opportunity to teach DARCI is open to the public with no restrictions through a persistent public website. What this amounts to, is that DARCI's interpretation of images is an unpredictable amalgamation of human sentiment and the machine learning algorithms at play behind the curtain. Effectively, DARCI presents a unique interpretation of the social culture surrounding her development—arguably the role of an artist.

Clearly, DARCI is limited by constraints far more restrictive than her human counterparts; at the same time, these constraints provide a unique viewpoint. During the study, we engaged in activities designed to explore the constraints of DARCI from an artistic perspective in an attempt to stimulate introspection about the process of creating art in the participants while helping us improve DARCI. The study culminated in *Fitness Function*, an interactive art show, hosted by Brigham Young University's Visual Arts Department, in which DARCI acted as the sole juror for any-and-all submissions made on site.

This paper will explore the insights gained from simulating current computationally creative systems with social activities. We will begin by detailing the processes by which DARCI currently learns and creates, and observe the constraints imposed by these processes. We will then describe the activities, designed by both artists and computer scientists, that were employed during the study to foster an understanding of DARCI's limitations and provoke discussions about creativity. We will illustrate how these activities sparked insights into the creative process by those participating in the study. We will also outline how these insights will be used to redesign DARCI and improve her capacity to be perceived as a creative agent. Finally, we will describe the art show, *Fitness Function*, and the response it received from those who participated in it.

DARCI

An important component of creativity is appreciation, or the ability to evaluate one's own created artifacts [?]. In order for DARCI to appreciate art, she must first acquire some basic understanding of art. For example, in order for DARCI to appreciate an image that is bright and cold, DARCI must first understand the concepts 'bright' and 'cold'. To do this, DARCI must learn to associate images with descriptions. We teach DARCI to associate images with descriptions by providing her with hundreds of images, each labeled with adjectives that describe it. We also provide negative labels,

or adjectives that the images are not. DARCI considers low level features of each image that deal with light, color, texture and shape. These image features are then associated with the provided adjectives through a collection of artificial neural networks [?]. The more examples of happy and non-happy images that are provided, the better DARCI is able to understand what it means for an image to be happy and be able to evaluate how happy an image is. We collect labeled images for DARCI through a website that allows people to label images with adjectives (<http://darci.cs.byu.edu>). The user is presented with a random image from DARCI's database and the user is asked to provide adjectives that describe that image. Additionally, there is a section of the website in which DARCI attempts to label an image with adjectives herself and the user can correct DARCI by tagging the adjectives that are incorrect, thus providing negative labels.

Once she has sufficiently learned enough adjectives, the next step for DARCI is to modify an image to match a list of adjectives. DARCI has a large set of image filters, similar to what would be found in software like Photoshop, which are capable of modifying images. DARCI uses genetic algorithms to explore different image filter combinations and to learn how to use these filters to modify images to convey the meaning of each adjective [?]. Genetic algorithms have often been described as a creative process due to their purposeful, yet unpredictable nature [?].

Genetic algorithms are governed by a subroutine called a *fitness function* that evaluates each artifact. The fitness function for DARCI generates a single score that measures how well a set of filters modifies an image to match the provided adjectives. This fitness function is the output of the artificial neural networks that have been learning image-to-adjective associations. For example, if DARCI were given a photograph and the adjectives 'bright' and 'happy', then DARCI would begin the creation process by trying out random sets of image filters on the photograph. For each set of image filters, DARCI would score how 'happy' and 'bright' the resulting image is using the artificial neural networks. The sets of filters with the lowest scores would be discarded, while DARCI would keep the sets of filters with the highest scores. DARCI would then alter some of these filters (*mutation*) and even try recombining the best sets of filters (*crossover*). In this way DARCI is continually searching for the set of filters that would modify the photograph to be the most 'happy' and 'bright'. This process continues for many iterations, after which DARCI applies the best set of image filters found and outputs the modified image. A high level overview of DARCI can be seen in Figure ??.

DARCI is a simplified model of a complex process and hence has several important limitations. For example, DARCI's language is limited to lists of adjectives. Her percepts are limited to low-level image features and her ability to modify images is limited to a finite set of image filters. DARCI has no social context or past experience from which to draw analogies and insights. Finally, DARCI's feedback consists of a single metric with no context. These limitations will be discussed in greater detail throughout the paper.




Figure 2. Overview of DARCI's artifact creation process.

ARTIST INTERPRETATIONS

We began the study by introducing the participants to DARCI. Soon after that, still early in the study, we asked several of the artists to come up with activities for the rest of the participants that they felt represented the processes underlying DARCI as they understood her. This was not only to solidify their understanding of DARCI, but also to help the artists begin to think about their own processes of creating art. The artists came up with activities that highlighted many of DARCI's aforementioned limitations. This section contains a description and analysis of three of the activities that the artists devised.

Copy of a Copy

Copy of a copy was based on the simple premise that as information is transferred, it degrades. The participants were divided into four groups. Each group created a black-and-white picture as a team using only pencil and ink. There were no other rules. Once each team had finished their pictures, they found a copy machine on campus and made a copy of the picture. They then made a copy of the copy. They proceeded to copy each copy for fifty iterations. The groups then convened to show their results and discuss what happened.

As expected, the quality of each copy deteriorated to the point that the original image could no longer be recognized. Some groups experimented with different patterns or shade strengths. While the fidelity of each copy decreased, interesting artifacts would sometimes occur due to the way the original images were drawn. The participants decided that from an artistic perspective the degradation, to a point, was not always a negative thing.

This activity was designed by an individual who was concerned that as DARCI was taught adjectives, the quality of her learning would be suspect. Said another way, since both DARCI's percepts and the communication between teacher and DARCI are not perfect, there will always be a degradation of information. This individual assumed that DARCI was meant to reflect the opinion of her teachers. After this activity, the general consensus was that this imperfection in learning was, in fact, a good thing and not unlike a human teacher-student interaction. The flaws in communication are part of what allow the novel to emerge. This idea was explored throughout the study and the many activities we participated in; but this activity set the tone.

Rotating Artists

This activity had participants creating collaborative drawings all using the same medium, pencil. An 18 x 24 inch sheet of paper was set up on a drawing board for each individual. Each sheet was associated with a random adjective. Then each person was assigned a random rule restricting the mechanics of their drawing. These rules included requirements such as "draw with your off hand", "draw from over the top of the paper", "draw with your teeth", etc. The participants were randomly assigned a sheet of paper to start on. They were instructed to begin a drawing that fit the adjective assigned to the sheet of paper. After five minutes, the participants rotated to another sheet of paper. They were instructed to continue the previous individual's drawing using the adjective connected to the sheet. Participants continued to rotate until everyone had contributed to every drawing. Individuals were required to use their restriction on every other rotation. Finally, when all of the drawings were completed, the participants were split into three groups and each group created a collage out of one third of the drawings.

This activity simulated DARCI on several levels. First, participants were given a task similar to DARCI's primary function: they were required to create an image to fit an adjective. Second, participants were given hardware restrictions, their drawing rule and use of only pencils. This was an attempt to constrain them in the way that DARCI is constrained to digital media using only filters. The restriction is not identical but it did force individuals to create without their complete faculties, giving them something to think about. Finally, there were two attempts to specifically simulate the way DARCI generates her artefacts. The first was in rotating people from drawing to drawing. The second was in cutting up the drawings and combining them in some form of collage. These activities were attempts to copy the crossover and mutation mechanisms used in the genetic algorithm driving DARCI's image creation. Mutation is an iterative mechanism not unlike the incremental alterations contributed to each drawing by new artists. Crossover does combine features from multiple solutions as the collages did. However, unlike actual genetic algorithms, there was no evaluation of the changes made at each step and no decision based on these evaluations as to how to proceed. Also, the goal of the activity differed from DARCI's in that the participants were creating a wholly original image rather than modifying a source.

This was the first activity of the study to directly simulate DARCI's processes. As such, it began to bridge the gap between the understanding of the artists and that of the computer scientists.

Systems

In this activity the designer attempted to more accurately simulate DARCI's genetic algorithm based on the discussion deconstructing *Rotating Artists*. In addition, the designer added some style by presenting the activity as if those participating were themselves machines. *Systems* was divided into two sessions called *System 1* and *System 2* respectively. Each participant was given a unique *Ink Hardware*, a tool for

SYSTEM 1, Hierarchy of Basic Alternatives

Final Image by Random System Purge

Instructions: Use the Ink Hardware you were provided to generate an image in the EVENT area. Use a pencil to input any necessary data in boxes below the event. The time allotment for each event is 1 minute mandatory mental planning + 15 minutes INK HARDWARE use.

INK HARDWARE: Randomly allocated ink-drawing tools. Rules are included with each tool set; each individual will be required to integrate the Rule associated with the tool for each EVENT.

EVENT#	PROCESSING UNIT ID:	PROGRAM:	SOURCE CODE:	SYSTEM PURGE:
3	23624MAD45245E	flowers	happy	eliminated if loser of thumb-war (held other's wrist, 3 count win)

EVENT: Area provided to generate an image using the allocated INK HARDWARE.

PROCESSING UNIT ID: Input prior to event// Initials + age + gender + Geographical location you define as home (definition may vary according to individual) + Ink Hardware Serial.

PROGRAM: Event instructions provided by DARCI proxy// Process according to your interpretation. Supplementary instructions are provided on the back.

SOURCE CODE: Input immediately following event// One word describing a prevailing or dominant quality of mind characterizing your time spent on this event (a.k.a. your temperament).

SYSTEM PURGE: After the time allotment for the EVENT is completed, pair up with another individual to initiate a system purge according to the given instructions.

Figure 3. The instructions for System 1 as presented to participants.

applying ink on paper (i.e. various pens, ink blots, an eye dropper, etc.). The participants were then presented with the instructions shown in Figure ?? and each given a card similar to the one in the instructions. Basically, within sixteen minutes, each participant created an image of the noun listed under *Program* using their “hardware”. They then wrote a word that described their experience under *Source Code*. Finally, they paired up with another random participant and performed the *System Purge* to determine which individual would progress in *System 1* and which would be eliminated and move on to *System 2*. Those that progressed within *System 1* got a new card with a *Program* to modify the eliminated participant’s art in some way. Individuals continued to pair up and *System Purge* until there was only one person left who then created the final image.

The results of the four rounds of elimination that occurred in *System 1* are shown in Figures ??-??(a) respectively. The final image is shown in Figure ??(b). This activity simulated the generational progression of DARCI’s genetic algorithm. Each generation, the “optimal” images would move on to compete in the next generation. Furthermore, each generation was built off of the previous one. While an analogue for crossover was not present, one for mutation was. The significant difference between *System 1* and DARCI’s genetic algorithm was that the criteria for success in an image was not directly correlated with the quality of the image. The winner of each *System Purge* was determined by either winning a thumb war, having the fewest siblings, being born furthest from the activity location, or wearing the most articles of clothing, depending on the round. Typically,




Figure 4. The results of the first round of System 1.




Figure 5. The results of the second round of System 1.

such an indirect evaluation criteria is not ideal in genetic algorithms. However, it can have unforeseen and fascinating implications (maybe people with strong thumbs make better artists).

As participants were eliminated, they moved on to *System 2* (see Figure ??). *System 2* was essentially a group project where each panel iterated on the previous one. People were assigned a panel to complete in the order they were eliminated from *System 1* and then by birth date for individuals eliminated simultaneously. Each panel was supposed to reflect the predetermined *Program* and the previous panel. Participants could use whatever *Ink Hardware* was available at the time they started their panel. This system explored the concept of evolution at a high level. Mutation occurred in the sense that individuals iterated upon one another’s work. The progress of the panels was aided by an evolutionary theme



Figure 6. The results of the third round of System 1.

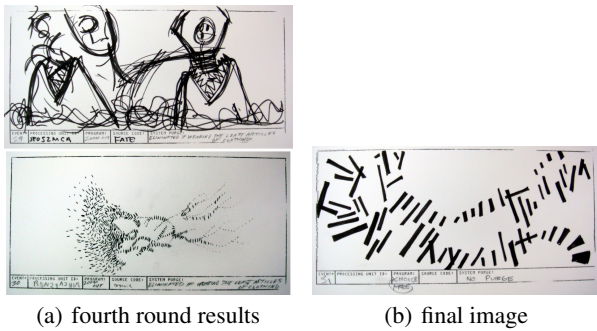


Figure 7. Fourth round and final image results for System 1.

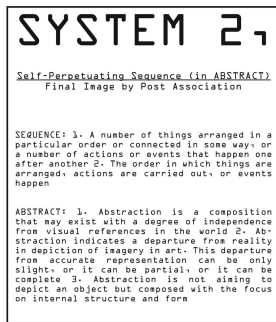


Figure 8. The instructions for System 2 as presented to participants.

in the *Program* titles. The *Program* themes progressed in the following order: birth, adaptation, bewilderment, fear of other, conflict, power, unification, stress, pain, relief, disbelief. The result of System 2 can be seen in Figure ?? . While open to interpretation, the *Program* themes hint at a message about the process of merging ostensibly conflicting viewpoints: say that of visual arts and that of computer science. With this in mind, *Systems* not only demonstrated an artistic interpretation of DARCI's creative process, but also acted as a metaphor for the study as a whole.

COMPUTER SCIENTIST INTERPRETATIONS

When an artist creates art, what is going on in their head? Said another way, given the artist's inputs, what determines the output? In this context, the input includes feelings, education, training, genetics, the weather, day-to-day senses—everything that one experiences. The output is the art. Some process occurs that transforms these inputs into artifacts. The activities that we designed were aimed at trying to understand this process. The first activity, *Bromerly*, was aimed at exploring the process of learning associations and expressing them in art. The second, *Celebrities Crossover*, was aimed at exploring the idea that meaning is not universal. The third activity, *External Evaluation*, was aimed at exploring the evaluation of art from both the creator's and the critic's perspectives.

Bromerly

When we, as humans, hear adjectives like 'happy', 'scary', or 'red', we have a lifetime of experiences and examples to



Figure 9. The results of System 2.

help us understand what these words mean. These adjectives fit into a cultural and environmental context that influences how we interpret them. DARCI has none of these advantages. When DARCI learns a new adjective like 'gloomy', she has no past experience to draw upon, and she has no cultural influences to help define it. All DARCI has are the example images of 'gloomy' and 'not gloomy' that are given to her.

An activity was designed for the human artists to simulate DARCI's processes and limitations as much as possible. The goal was to see how they would compare to DARCI, to determine what their creative processes would be like in a restricted setting, and to discover improvements for DARCI. Three fake adjectives were invented that have no reference to any other word, and have no background or cultural meaning. These adjectives are 'orfly', 'bromerly', and 'flamping'. Unbeknownst to the human artists in the study, these adjectives actually represented a real adjective that DARCI knew fairly well. Orfly meant weird, bromerly meant blocky, and flamping meant lonely. The artists were then provided with the same example images for each adjective that DARCI was trained on (both positive and negative examples).

The participating artists were split into three groups. Each group was given one of the fake adjectives and the example images for that adjective. The artists spent several minutes looking over the provided images to learn as best they could what the fake adjectives meant. Each artist was then given their own source image and instructed to create something based off the source image that reflected the adjective that they just learned. Essentially they had to communicate the meaning of their adjective through their creation. Some of the resulting images along with their sources can be seen in Figure ?? . After everyone was done, all the paintings




Figure 10. Three example images (with their sources) created during the Bromerly activity. One example of each fake adjective is presented. Flamping means lonely, orfly means weird, and bromerly means blocky.



Figure 11. The cluster of flamping (lonely) images.

were put together. A random group of college students, not involved with the study, was then brought in to divide the paintings into three groups. These people had no knowledge of the activity or how they were supposed to group the paintings. The volunteers were able to group the paintings of the same adjectives together with the exception of one painting. The grouping results can be seen in Figures ??-??.

Although unable to specifically define their adjective, the artists were able to identify certain image features or qualities that they associated with their adjective. For example, they were able to tell that bromerly (blocky) images tended to have hard geometric lines, tended to be very structural, and tended to have a lot of divided space. They could tell that orfly (weird) images often had exaggerated curvy shapes and had some distinct peculiarity. They could also tell that flamping (lonely) images were usually dark and had a singular focus. The artists were then able to incorporate these characteristics into their own paintings. During the grouping phase, the people doing the grouping were able to pick up on these features and successfully group similar paintings together, while only miss-classifying one painting. This notion of evaluating the paintings by grouping led to a paper currently in submission that explores the use of clustering algorithms to evaluate DARCI’s artifacts.



Figure 12. The cluster of orfly (weird) images. The far left image should have been clustered with bromerly (blocky).




Figure 13. The cluster of bromerly (blocky) images.

A notable difference between DARCI and the human artists was how quickly humans were able to learn the adjective with so few example images. This is not surprising considering how sophisticated human vision is. DARCI needs hundreds of example images and can only see explicitly specified low level image features, while the human artists were able to quickly identify and extract the useful image features and then convey those features into their own paintings. Taking note of the way the human artists examined the example images and identified the important image features has provided us with useful insights into the type of image features DARCI needs to extract.

Through this activity, the artists gained a better understanding of how DARCI learns adjectives with the limitations that she has. The artists found it very interesting to learn a new word with only images rather than linguistically. They mentioned that when most words are heard, an associated image usually comes into people’s minds. For example, if someone hears the word ‘chair’, an image of a chair pops into their mind. It may be a different type of chair for each person, but some kind of image pops into their head.

“When I say chair, everybody thinks of a chair in their mind. They see that image, it might be a different chair. I just think that the images are the words.”

They also said that using visuals to convey the mean of a word is less ambiguous than using other words to explain the meaning because those other words have different meaning and connotations for different people, while images are more concrete.

“It is easier for us to get it visually than through words, words are very ambiguous. ... I usually have to see something visually to remember it.”

In order to provide a more direct comparison with DARCI, several artists were asked to learn additional fake adjectives and then, using only image filters found in Photoshop, modify a source image to reflect multiple adjectives. Restricting the artists to Photoshop ensured that they were limited to the same tools that DARCI has to modify an image. We had DARCI do the same task with the same adjectives and same source images. The results can be seen in Figures ??-??. The artists were again able to quickly identify the distinctive image features that they associated with their adjectives. Some artists were unfamiliar with Photoshop and spent time trying different image filters to learn how they would affect the source image. They would then find the ones that modified the source image to best reflect their understanding of




Figure 14. Images produced by a human and DARCI from the same source image for the combination of the adjectives *bromerly* (blocky), *blemistic* (majestic), and *diambic* (avian).




Figure 15. Images produced by a human and DARCI from the same source image for the combination of the adjectives *flamping* (lonely) and *orfly* (weird).

the adjectives. This process that they described is very similar to the genetic algorithm process DARCI uses to learn images filters to modify images. Additionally, these results are consistent with prior research that shows that DARCI is capable of producing images comparable to human artists under similar constraints [?].

Celebrities Crossover

This activity was structured around the party game *Celebrities*. In this game, players secretly write the names of several people or characters (fictional or not, famous or not, made up on the spot or not) on about 40 pieces of paper. The players are divided into two teams. Each team takes 30 second turns by having one of its members describe the characters on the pieces of paper while the other members try to guess the character. Every correct guess scores the team a point. This continues until all of the characters have been guessed. This process is repeated for two more rounds using the same set of characters and same teams but with different rules. During the second round, players can only use a single word to describe the character. On the third and final round, players can only use gestures to describe the characters. Since the same characters are used every round, as the game progresses, players develop unique associations to aid in describing these characters. Often these associations have nothing to do with a character at all, but stem from players attempting to describe characters they don't know as fast as



Figure 16. Images produced by a human and DARCI from the same source image for the combination of the adjectives *orfly* (weird), *blemistic* (majestic), and *diambic* (avian).

possible. These frequently bizarre associations get picked up by everyone in the game (both teams) and can even become memes within the group that played the game together.

The study participants were divided into two groups (with two teams in each group). Each group played *Celebrities* with the exact same character names, but physically separated from each other. Thus, two divergent sets of character associations were allowed to develop. Once the groups had completed their games, one team from each group was transplanted into the other group. Each of the resulting groups then played a revised version of *Celebrities*. The new version of the game was no longer a team game; instead, it was everyone for themselves. The participant with the most points from each group would be the winner. The rules for the game were the same as round two from *Celebrities* (meaning that participants could only say a single word or use gestures) except that each participant was performing for everyone else in the group instead of teammates. Each individual that correctly guessed a character was awarded a point; the person performing only received points if the player who correctly guessed the character came from a different group than their original one. If this happened, the performing player was awarded 3 points. This forced the participants to deliberately avoid the associations they had formed previously, while developing new ones that would give everyone playing the same chance of guessing correctly.

In semiotics, the association between a sign and that to which the sign refers denotes meaning. An important aspect of visual arts is the ability of the artist to convey meaning to an audience through their art. Thus, the image that an artist produces would be the sign that is conveying the meaning that the artist intends. We designed this activity to stress the idea that not every culture associates signs with the same meanings. In this context, the term culture can refer to social groups of any scope, including an individual. When an artist creates a piece of work, to what extent is the artist communicating meaning? Is the artist only communicating to those that share the same associations as intended in an artifact? Do artists make an effort to communicate outside their personal semiotic domains [?]? We intended for this activity to force those participating to ask themselves these questions. In answering these questions, we hoped to learn the extent to which the communication of meaning would need to be emulated in DARCI for her to be accepted as an artist.

In *Celebrities*, the signs are the gestures or single words, and the meanings are the character names. With DARCI, the signs are the feature vectors that she extracts from an image, and the meanings are adjectives. Since the feature vectors represent DARCI's percepts of an image, we suggest that DARCI's signs are effectively the images themselves. In *Celebrities Crossover*, each group forms its own semiotic domain for the set of character names. When the groups cross-pollinate, they are required to communicate outside these newly formed semiotic domains. In our experience, this involved relying upon alternative associations that existed prior to the activity—in other words experience outside the activity's restricted gesture/word based symbolism.

One shortcoming of DARCI that was brought to our attention by these observations was that currently, DARCI has no associations besides the image-adjective relationship. This means that she is unable to use experiences external to these associations to communicate with an audience outside her personal semiotic domain. Another shortcoming we discovered was that, even within the image-adjective relationships, DARCI is only able to model a single semiotic domain. This domain represents a function of all of the opinions of people who have trained DARCI weighted by how much time they have spent with her. While the associations this creates may be interesting, there is no attempt to cater to a specific audience. We are redesigning DARCI so that she will be able to model each individual that trains her in addition to a general model. This will allow her to cater to specific individuals when commissioned—essentially communicating within that person’s semiotic domain for image-adjective associations. Whether DARCI should cater to individuals is another question; but, at least having the capability to communicate more effectively is something we are interested in.

Referring to *Celebrities Crossover*, the following comment was made by a participating artist:

“I was so fascinated by the associations that we made. Dan said he didn’t think it was a big deal that we made “incorrect” (insane, irrational...) associations, but does DARCI make associations that are incorrect? Does anything she do come out of random associations or is all of her information supplied by us? What happens if the associations she makes are incorrect? Maybe nothing. Maybe it doesn’t matter.”

Another artist replied with:

“The associations being created are not always expected. They may not be the prophesied outcomes. But that does not make them incorrect. Each association has been built upon an individuals [sic] experiences. So for that individual the association is alive and correct. Even if it is not the association that was predicted.”

This discussion illustrates some ideas that were brought to the participant’s attention as they pondered the activity. Whether or not these ideas are actually valid in the domain of visual arts is irrelevant. What’s important is that the participants are thinking about the process of creation—their process of creation.

Abstracted Feedback

As described previously, DARCI uses genetic algorithms to learn what set of image filters will modify an image to reflect a certain adjective. This genetic algorithm uses a fitness function to guide the exploration. For example, if DARCI was trying to make an image fiery, then DARCI would try several sets of image filters. The modified images would then each get a numerical score from the fitness function. This score tells DARCI how fiery each image is. This is just a single number—it does not tell DARCI anything about how or why it is or isn’t fiery. All DARCI knows is that the




Figure 17. The template used to evaluate the objective group.




Figure 18. Examples of drawings created by the group evaluated subjectively based on how scary their images were. Drawings are chronologically ordered from top left to bottom right.

higher the number the better. DARCI can then take the highest scoring sets of image filters and perform crossover and mutation on them to search for even higher scoring sets of image filters.

An activity was designed to give human artists this same limitation in feedback in order to see how they would perform, and to see what their thought processes would be. Each person was given a quarter sheet of paper and had to draw a picture that matched a specific criterion. To start with, no one had any idea what that criterion was. The only direction they were given was that their drawing would be scored based on some criterion unknown to them and they had to try to maximize their score. After each person submitted a drawing, the drawings were each given a numerical score. The drawings were handed back with their score, and each person could view the returned drawings and try to figure out why one drawing scored higher than another. Each person then submitted another drawing in order to increase their score. This process was repeated for several iterations.

The participating artists were split into two groups each with different criteria. The first group had a purposefully subjective criterion for which they had to draw a scary image. The drawings were each given a score based subjectively on how scary the human evaluators thought the drawing was. The second group’s criterion was more objective. They had to draw a specific shape with the correct size and location on the paper. Human evaluators scored each drawing on how well it matched a specific template (seen in Figure ??) by holding the template up to the drawing and counting how many square regions on the template overlapped with the drawing. Example drawings from group 1 and group 2 can be seen in Figures ?? and ?? respectively.




Figure 19. Examples of drawings created by the group evaluated objectively based on how closely their images fit a grid template. Drawings are chronologically ordered from top left to bottom right.

The artists got frustrated quickly when they would put time and effort into a drawing and then receive a low score and not know why. The first group at first submitted a very wide range of drawings. Eventually, someone drew a pirate ship which bumped up their score considerably. The others in the group keyed in on this higher score and all started drawing pirate related things. Soon every artist in group 1 converged to drawings of skulls. The second group was never able to get the shape completely right, they converged to a drawing of a shaded square. They did, however, converge roughly to the correct size and location on the paper.

This activity proved to be difficult and frustrating for the artists. To successfully submit a high-scoring drawing required a lot of time, persistence, and patience. This activity helped the artists understand how DARCI has to deal with limited feedback. It quickly generated a discussion on how art is evaluated in the real world. Many of the artists who had submitted their artwork to various art exhibits admitted that they were usually oblivious to the requirements of a particular art show. They would just submit all of their best work and hope something got in. This discussion developed into the idea for the art exhibit *Fitness Function*.

FITNESS FUNCTION

Fitness Function is an interactive art exhibit where DARCI acts as the sole juror for pieces submitted throughout the show. The exhibit begins with an empty wall that gets filled with art that DARCI selects from submitted work. DARCI is set up in the exhibit with a simple interface for submitting images via USB sticks. Participating patrons provide their name, email, the title of the piece, and the image itself and then DARCI scores the image between 0 and 100 using the neural networks mentioned previously. Images that receive a score of 70 or higher are printed for the artist to hang on the wall, and images that receive a score of 90 or higher receive a jury award. Artists can submit as many images as they want and anyone is eligible.

With the collaboration of Brigham Young University's Visual Arts department, we set up *Fitness Function* in the Harris Fine Arts Center to operate between March 19th and 30th, 2010, with a closing reception on the 31st. Advertisements for the exhibit were spread throughout campus via school website, posters, fliers, and word of mouth.



Figure 20. The ten images that received jury awards at *Fitness Function*.

We used a model of DARCI trained exclusively by the participants in our study. Just as potential images created by DARCI's own genetic algorithms are evaluated based on how closely they match particular adjectives (described earlier), the images submitted by participants of *Fitness Function* were evaluated based on how closely their images matched a set of secret adjectives that we chose for the exhibit. The adjectives we chose were 'phantasmagoric' and a special artificial adjective that we created for the purpose of *Fitness Function*. The special adjective was DARCI's prediction of whether a piece would be accepted to Brigham Young University's Museum of Art. In other words, the adjective was 'museum-of-artsy'. So that the artists involved in our study could participate, the adjective 'phantasmagoric' was kept secret from them; they weren't even told about the existence of 'museum-of-artsy'.

By the end of the exhibit, 93 people had submitted 787 images. 145 of these images were accepted, and 10 of those received a jury award at the closing reception. The jury award winners are shown in Figure ???. A website documenting the show can be found at <http://darci.cs.byu.edu/fitnessfunction/>.

As indicated by the comment book in the exhibit, word of mouth, personal observations, and the nature of submissions, people's reactions to the exhibit varied significantly. Some felt pride and vindication due to the acceptance of their art: "I MADE IT IN! YEAH!". Some felt frustration: "Fun and frustrating at the same time. More like a game to try and figure out than serious." Others took it as a personal challenge. One individual submitted 106 images, clearly experimenting with lighting, image quality, and other less conventional alterations (Figure ??). Thirty-eight of his images were accepted. As the last entry in the comment book he wrote, "IT WAS FUN WHILE IT LASTED DARCI." Whatever their opinion, people couldn't help but personify DARCI: "An interesting project, however, it seems that DARCI is very opinionated in her ratings. From a look at the wall, she has some rather odd tastes."

In this exhibit DARCI represents a metaphor for all kinds of external evaluation criteria to which an artist is subject, whether implicitly, explicitly, consciously, or unconsciously. We wanted to challenge people's notions of evaluation. How does one decide that they have succeeded with a piece of art? Does one only esteem a work if it is valued by others; if it can get into an art show? What if one's art gets a high score from a machine? That was quite significant for some peo-



Figure 21. Two sample submissions from a witty artist.

ple. The ability to evaluate artifacts on some level is critical in any act of creativity, and from this exhibit we learned that people at least perceive that DARCI has an opinion (i.e. that DARCI is capable of evaluation). The perception of creativity has been argued to *be* creativity [?]. If this is the case, then shows like *Fitness Function* are a step in the right direction for indentifying arguably creative systems.

CONCLUSIONS

During the study, the question was raised: “does DARCI think?” The computer scientists responded with “if we can describe it, we can program a computer to do it.” This quote encapsulates the computer science viewpoint and our motivation for the study. If we can formalize the process of thinking, then a computer can be programmed to think. If we can formalize what it means to be creative, then a computer can be programmed to be creative. The activities in the study were designed to elicit personal introspection on one’s creative process in the hopes of shining light on that elusive process. By focusing the activities on DARCI’s limitations, we were able to help those participating in the study isolate the thoughts and actions behind their creations.

As illustrated in the previous sections, we were able to gain some insight for our system, DARCI, through the activities we designed and the discussions that followed: the degradation of information associated with learning is acceptable, clustering algorithms can be used in the evaluation of artifacts, more realistic image feature extraction is vital, the creation of perceptual associations is important (perhaps in all creative systems), genetic algorithms do appear to emu-

late part of the creative process, the creation of unique relationships is acceptable and useful, being able to create relationships across multiple senses is necessary to communicate across semiotic domains, and finally the external evaluation of one’s work is significant in the creation process.

Throughout the study, participants were encouraged to contribute to a blog titled Art[ificial] (<http://computationalcreativity.blogspot.com/>). This blog was designed for participants to leave their thoughts on the activities and discussions we engaged in during the study. Consider the following entry:

“Today I was thinking of interpretation. I interpret a lot of things. I think that’s a good word to put to my process of taking a lot of input and making an output that I feel will have significance in a realm beyond my personal self. I am interpreting emotions, thought processes/concepts, words, into a visual language. Other words I think might be synonymous are ‘translating,’ or ‘encoding.’ Often times I have heard the word ‘feel’ used in critiques, or have thought it to myself, where something simply feels consistent, or right, or something feels like it doesn’t belong. It is a way of interpreting things across different areas, from visual to emotional, or visual to psychological—a synesthesia of emotion, subconscious, and physical senses.

I read some interesting articles about synesthesia and creative processes, and I think that is something significant to art making—the ability to have ideas cross over from one realm to another, or one language to another, while still maintaining a certain degree of integrity or self sufficiency.”

This thought sums up the nature and purpose of the study—the synthesis of artistic and scientific interpretations for the goal of understanding the process of creation.

ACKNOWLEDGMENTS

We thank Professor Joseph Ostraff and his VASTU 480/680 students for the fruitful collaboration.

This material is based upon work supported by the National Science Foundation under Grant No. IIS-0856089.