Using a Contextual Focus Model for an Automatic Creativity Algorithm to Generate Art Work

Steve DiPaola
Simon Fraser University, Vancouver, CA
sdipaola@sfu.ca

Abstract

We sought to implement and determine whether incorporating cognitive based contextual focus into a genetic programming fitness function would play a crucial role in enabling the computer system to generate art that humans find "creative" (i.e. possessing qualities of novelty and aesthetic value typically ascribed to the output of a creative artistic process). We implemented contextual focus in the evolutionary art algorithm by giving the program the capacity to vary its level of fluidity and functional triggered dynamic control over different phases of the creative process. The domain of portrait painting was chosen because it requires both focused attention (analytical thought) to accomplish the primary goal of creating portrait sitter resemblance as well as defocused attention (associative thought) to creativity deviate from resemblance i.e., to meet the broad and often conflicting criteria of aesthetic art. Since judging creative art is subjective, rather than use quantitative analysis, a representative subset of the automatically produced art-work from this system was selected and submitted to many peer reviewed and commissioned art shows, thereby allowing it to be judged positively or negatively as creative by human art curators, reviewers and the art gallery going public.

Keywords: Evolutionary Systems, Genetic Programming, Contextual Focus, Creativity, Computational Modelling

1 Introduction

Creativity is a complex set of cognitive process theorized to involve, among other elements, attention shifts between associative and analytical focus (Gabora, 2000), novel goals (Luo and Knoblich, 2007), and situated actions and difficult definitions of evaluation. Computational creative systems (CES) strive to model a variety of creativity's aspects using computer algorithms from evolutionary 'small-step' modifications to intelligent autonomous composition and 'big-leap' innovation in an effort to better understand and replicate creative process (Boden, 2003). The focus by some researchers on replicating creativity in computational algorithms has been instrumental in learning more about human cognition (individual and collaborative) and how creative support tools might be used to enhance and augment human creative individuals and teams. All these aspects continue to evolve our perceptions of creativity and its role in computation in the current technology-saturated world.

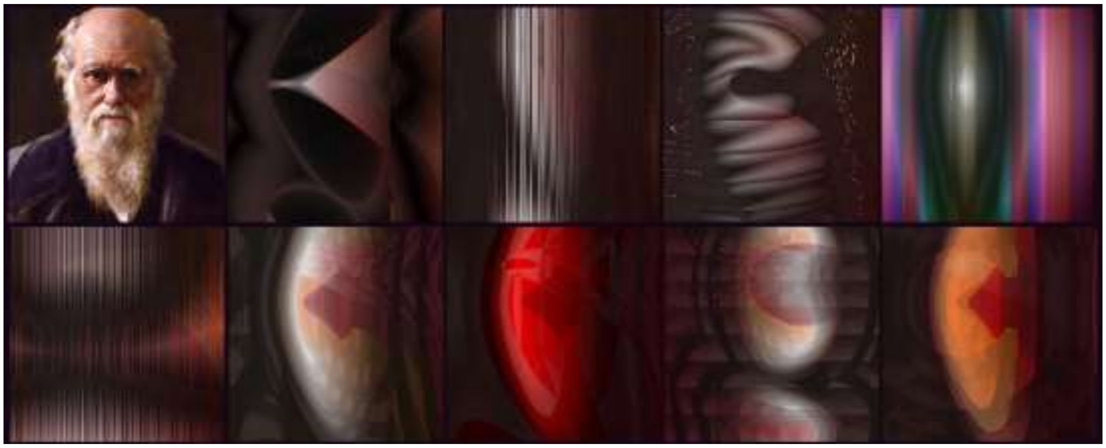

Figure 1. Source Darwin image (top left) with examples of our evolved abstract portraits created using our DarwinsGaze autonomous creative genetic programming system.

Systems modeling creativity computationally have gained acceptance in the last two decades, situated mainly as artistic and research projects. Several researchers in computational creativity have addressed questions around such computational modeling by outlining different dimensions of creativity and proposing schema for evaluating a "level of creativity" of a given system, for example (Ritchie, 2007; Jennings, 2010; Colton, Pease and Charnley, 2011). While there is ongoing research and scholarly discourse about how a system is realized, how the results are generated, selected and adjusted and how the process and product are evaluated, there is less research about direct applications of creative cognitive support systems in real-world situations.

2 Contextual Focus: Associative and Analytical Thinking

We explore creativity from theories of cognition that attempt to understand attentional shifts between associative and analytical focus – what we call “contextual focus” or “contextual fluidity”. The existence of two stages of the creative process is consistent with the widely held view that there are two distinct forms of thought (Neisser, 1963; Piaget, 1926; Sloman, 1996). It has been proposed that creativity involves the ability to vary the degree of conceptual fluidity in response to the demands of any given phase of the creative process (Gabora, 2000, 2002; DiPaola & Gabora, 2009). Again, this dimension of variability in focus is referred to as contextual focus. Focused attention produces analytic thought, which is conducive to manipulating symbolic primitives and deducing laws of cause and effect, while defocused attention produces fluid or associative thought which is conducive to analogy and unearthing relationships of correlation. Thus, creativity is not just a matter of eliminating rules but of assimilating and then breaking free of them where warranted. Said another way, divergent or associative processes are hypothesized to occur during idea generation, while convergent or analytic processes predominate during the refinement, implementation, and testing of an idea. This is referred to as contextual focus because it requires the ability to focus or defocus attention in response to the context or situation one is in. Defocused attention, by diffusely activating a broad region of memory, is conducive to divergent thought; it enables obscure (but potentially relevant) aspects of the situation thus come into play. Focused attention is conducive to convergent thought; memory activation is constrained enough to hone in and perform logical mental operations on the most clearly relevant aspects.

This paper focuses on the implementation and applicability of contextual focus through our research system, DarwinsGaze, developed to use an cognitive based automatic fitness function inspired by contextual focus human creativity. Our analysis of their process combined with our knowledge of the cognitive aspects of creativity (gleaned from our early research), was used to design and implement the DarwinsGaze system. We concentrate on the qualitative impact made by the explicit incorporation of contextual focus into the system as a whole, and its ability to elevate the perceived quality and novelty of system output to a level audiences judged reminiscent of successful "artistic, human-style" creativity.

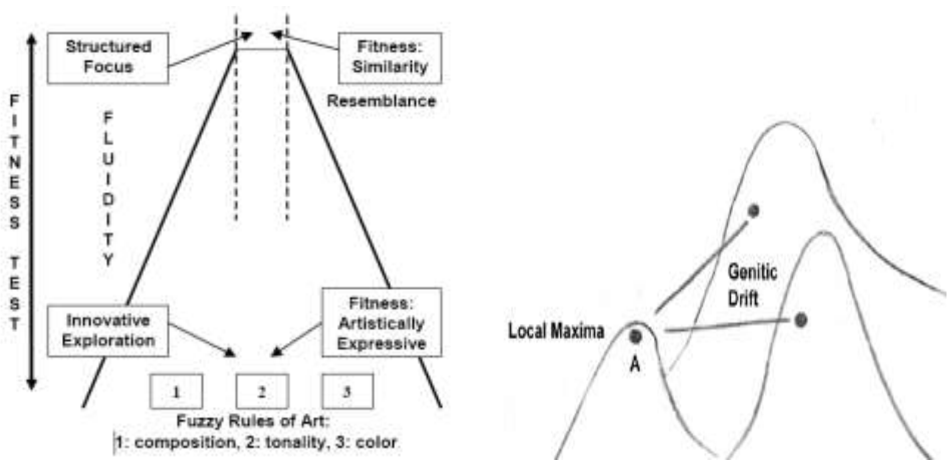

Figure 2. We model the cognitive architecture of contextual focus (left): resemblance simulates analytical thinking, 3 fuzzy art rules simulate associative thinking, functionally triggering between them. On a local maxima, the system automatically goes wide in the search space (right) in an Aha moment or back to refinement.

3 The DarwinsGaze System

The DarwinsGaze system (DiPaola and Gabora, 2009) is a Creative Evolutionary System (CES) (Bentley and Corne, 2002) (see Figure 3) based on a variant of Genetic Programming (GP). Unlike typical Genetic Programming systems this system favors exploration over optimization, finding innovative or novel solutions over a preconceived notion of a specific optimal solution. It uses an cognitive theory based automatic fitness function (albeit one specific to portrait painting) allowing it to function without human intervention between being launched and obtaining the final, often unanticipated and pleasing set of results; in this specific and limited sense we refer to DarwinsGaze as "autonomous". The inspiration for this work is to directly explore to what extent computer algorithms can be creative on their own (Gabora and DiPaola, 2012). Related work has begun to use creative evolutionary systems with automatic fitness functions in design and music (Bentley, 2002), as well as building of a creative invention machine (Koza et al, 2003). Typically these systems allow a human user to pick those individuals that will be mated – making the human the creative judge. In contrast, our system used a function trigger mechanism within the contextual focus fitness function which allowed the process to run automatically, without any human intervention once the process was started. It was not until the evolutionary art process came to completion that humans looked at and evaluated the art. So the contribution of the DarwinsGaze work is to model, in software, newly theorized aspects of cognitive based human creativity, especially in terms of fluid contextual focus (see Figure 2).

DarwinsGaze capitalizes on recent developments in GP called Cartesian Genetic Programming (CGP) (Miller 2010). CGP uses GP techniques (crossover, mutation, and survival), but differs in certain key respects. The program is represented by a directed graph of indexed nodes. Each node has a number of inputs and a function that gives an output based on the inputs. The genotype is a list of integers determining the connectivity and functionality of the nodes, which can be mutated and mated to create new directed graphs. CGP has several features that foster creativity including 1) its node based structure facilitates the creation of visual mapping modules, 2) its structure can represent complex computational input/output connectivity, thus accommodating our sophisticated tone and temperature-based color space model which enables designerly decision making, and most importantly 3) its component-based approach favors exploration over optimization by allowing different genotypes to map to the same phenotype. The last technique uses redundancy at the input, node, and functional levels, allowing the genotype to contain nodes that are not connected to the output nodes and so not expressed in the phenotype. Having different genotypes (recipes) map to the same phenotype (output) provides CGP with greater neutrality (Yu and Miller, 2005). Our work is based on Ashmore and Miller's (2004) CGP application to evolve visual algorithms for enhanced image complexity or circular objects in an image. Most of their efforts involve initializing a population and then letting the user take over. Our system was based upon their approach, but was significantly expanded with a more sophisticated cognitive based (contextual focus) creativity function, and revised the system for a portrait painter process.

Our GP function set has 15 functions (see Figure 3) which use unitized x and y positions of the portrait image as variables and additional parameter variables (noted PM) that can be affected by adaptive mutation. Functions are low level in nature which aids in a large 'creative' search space, and output HSV color space values between 0 and 255. An individual in our population is manifested as one program that runs successively for every pixel in the output image, which is then tested against our creative fitness function. This allows correlated painterly effects as one moves through the image. Functions 1 - 5 use simple logical or arithmetic manipulations of the positions (low level functions create a larger creative search space), whereas 7 - 14 use trigonometric or logical functions that are more related to geometric shapes and color graduations.

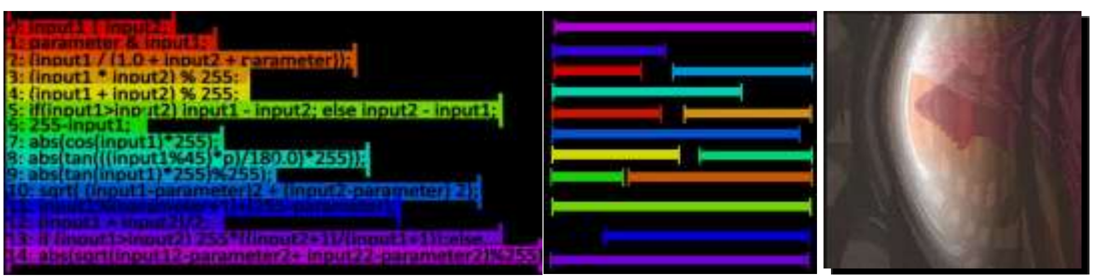

Figure 3. Our modified biologically inspired CGP system which uses crossover, mutation and replication in a contextual focus based fitness function, the 15 genes (left) produce Java programs that generate a visual output (right) in populations, testing by our cognitive inspired fitness function which ran without human intervention.

The contextual focus based fitness function varies fluidly from tightly focusing on resemblance (similarity to the sitter image, which in this case is an image of Charles Darwin), to swinging (based on functional triggers) toward a more associative process of the intertwining, and at times contradicting, 'rules' of abstract portrait painting. Different genotypes map to the same phenotype. This allows us to vary the degree of creative fluidity because it offers the capacity to move though the search space via genotype (small ordered movement) or phenotype (large movement but still related).

For example, in one set of experiments this is implemented as follows: if the fittest individual of a population is identical to an individual in the previous generation for more than three iterations, meaning the algorithm is stuck in analytic mode (local maxima) and needs to open up, other genotypes that map to this same phenotype are chosen over the current non-progressing genotype, allowing divergent open movement through the landscape of possibilities.

The automatic fitness function partly uses a ‘portrait to sitter’ resemblance. Since the advent of photography, portrait painting has not just been about accurate reproduction, but also about using modern painterly goals to achieve a creative representation of the sitter. We have created a fitness function that mainly rewards accurate representation, but given certain situations it also dynamically functional triggers to reward visual painterly aesthetics using three simple rules of art creation as well as a portrait knowledge space. Specifically, the aesthetic portion of our fitness function 1) weighs for face (centered) versus background composition, 2) uses tonal similarity over exact color similarity matched with a sophisticated artistic color space model which weighs for warm-cool color temperature relationships based analogous and complementary color harmony rules and 3) employs unequal dominate and subdominant tone and color rules and other artistic rules based on a portrait painter knowledge domain (DiPaola and Gabora, 2009) as illustrated in Figure 2. We mostly weight heavily towards resemblance to start, which gives us a structured system, but can under the influence of functional triggers allow for artistic creativity. In this way fitness function scores of resemblance simulates the analytical thinking of working specifically to resemblance in the fitness function, while our three fuzzy rules of art simulate associative thinking, with sophisticated functional triggering built into our architecture between both in a way that simulates human contextual fluidity. The approach gives us novelty and innovation from within, or better said, responding to a structured system -- a trait of human creative individuals.

Generated portrait programs in the beginning of the run will look less like the sitter but from an aesthetic point of view might be highly desirable, since the function set has been built with painterly rules. Specifically, the fitness function in the DarwinsGaze system calculates four scores (resemblance and the three fuzzy associative painterly rules) separately and fluidly combines them in different ways to mimic cognitive based human creativity by moving between restrained focus (analytical resemblance) to more unstructured associative focus (3 rules of composition, tonality and color). To move fluidly between the two cognitive modes, in its default state the fitness function uses a ratio of 80% resemblance to 20% non-proportional scoring of our three painterly rules. Several functional triggers can alter this ratio in different ways to simulate contextual focus. The system will also allow very high scoring of painterly rule individuals to be accepted into the next population. When a plateau or local minima is reached for a certain number of epochs, the fitness function ratio switches course where painterly rules are weighted higher than resemblance (on a sliding scale) and work in conjunction with redundancy at the input, node, and functional levels.

Similarly, but now in reverse, to the default resemblance situation, high scoring resemblance individuals can pass into the next population when a percentage of painterly rule individuals is met. Using this more associative mode, high resemblance individuals are always part of the mix, and when these individuals show a marked improvement, a trigger is set to return to the more focused 80/20 resemblance ratio.

As the fitness score increases, portraits look more like the sitter. This gives us a somewhat known spread from very primitive (abstract) all the way through to realistic portraits. Thus in effect the system has two ongoing processes: (1) those most ‘fit’ portraits that pass on their portrait resemblance strategies, making for more and more realistic portraits—the family ‘resemblance’ patriarchs, and (2) the creative ‘strange uncles’: related to the current ‘resemblance fit’, but portraits that are more artistically creative or artistically fit. This dual evolving technique of patriarchs and ‘strange uncles’ mimics the interplay between freedom and constraint that is so central to creativity. Paradoxically, novelty often benefits from the existence of a known framework reference system to rebel and innovate from. Creative people use some strong structural rules (as in the templates of a sonnet, tragedy, or in this case, a resemblance to the sitter image) as a resource or base to elaborate new variants beyond that structure (in this case, an abstracted variation of the sitter image).

Figure 4. Output from our automatic cognitive based CES system have been seen by thousands in the last few years and have been perceived as creative art works on their own by the art public, including 6 commission or juried major art shows including at MIT Museum (top) and the Cambridge University Kings Art Centre (bottom).

4 Results & Conclusion

The automatic creative output was generated over thirty days of continuous, un-supervised computer use. The images in Figure 4 show a selection of representative portraits produced by the system. While the overall population improves at resembling Darwin‘s portrait, what is more interesting to us is the variety of recurring, emergent and merged creative strategies that evolve as the programs in different ways to become better abstract portraitists (see www.darwinsgaze.com).

Humans rated the portraits produced by this version of portrait painting program with contextual focus as much more creative and interesting than a previous version that did not use contextual focus, and unlike its predecessor, the output of this program generated public attention worldwide. Example pieces were framed and submitted to galleries as a related set of work. Care was taken by the author to select representational images of the evolved un-supervised process; however creative human bias obvious exists in the representational editing process. Output has been accepted and exhibited at six major galleries and museums including the TenderPixel Gallery in London, Emily Carr Galley in Vancouver, and Kings Art Centre at Cambridge University as well as the MIT Museum, and the High Museum in Atlanta, all either peer reviewed, juried or commissioned shows from institutions that typically only accept human art work. A typical gallery installation consisted of 40-70 related portraits produced in time order over a given run. Gallery showings focus on "best resemblances" and those that are artistically compelling from an abstract portrait perspective. This gallery of work has been seen by tens of thousands of viewers who have commented that they see the artwork as an aesthetic piece that 'ebb and flows through seemly creative ideas' even though they were solely created by an evolutionary art computer program using contextual focus. Note that no attempt to create a pure 'creativity Turning Test' was attempted. Besides the issues surrounding the validity of such a test (Pease and Colton, 2011), it was not feasible in such reputable and large art venues. However most of the thousands of causal viewers assumed they were looking at human created art. The work was also selected for its aesthetic value to accompany an opinion piece in the journal Nature (Padian, 2008), and was given a strong critical review by the Harvard humanities critic, Browne (2009). While these are subjective measures, they are standard in the art world. The fact that the computer program produced novel creative artifacts, both as single art pieces and as a gallery collection of pieces with interrelated themes, using contextual focus as a key element of its functioning, is compelling evidence of the effectiveness of contextual focus.

5 Future Work

Many significant research CES systems exist that are both innovative and useful, some with strong cognitive architectures. However as the field matures, there will be an increasing need to make cognitive based CESs production worthy and work within a creative industry environment such as a digital design firm. To support others in this effort for production-targeted transformation, we have become to shift from an autonomous fitness function based creative system, DarwinsGaze, to an interactive fitness function based creative support system, Evolver (DiPaola et al, 2013), for real-world design collaboration. DarwinsGaze operates using a complex automatic fitness function to model the cognitive based theory of contextual focus as well as other aspects of human creativity simulated internally. In shifting to the Evolver project we found that the contextual focus perspective remained relevant, but now re-situated to overlay the collaborative process between designer and system. Four design principles developed on this basis were: 1) support analytic focus by providing tools tailored to the designer's specific needs and aesthetic preferences, 2) support associative or intuitive focus by relieving the designer's cognitive capacity, enabling a quick and serendipitous workflow when desired, and offering a large variety of parameterized options to utilize, 3) support a triggering of focus-shift between the designer and the system through options to 'bookmark' and save interesting pieces for later, as well as to move creative material from and to the system while retaining the work's semantic structure and editability, and 4) support a joint 'train of thought' between system and user by structuring a genotype representation compatible with human visual/cognitive intuition.

We found that the shift to a real-world design scenario required attention to the collaboration and creative processes of the designers who value their experience-developed expertise. The system design had to act as both a support tool engaging some cognitive load of the process, and a flexible, interactive repository of potentially successful options. Future real-world design considerations can explore methods for adapting intelligent operations to the cognitive processes and constraints of necessary situations, taking into account the expertise of collaborators.

6 Acknowledgements

This work was supported by Liane Gabora, Graeme McCaig, Lawrence Ashmore and Julian Miller and that of the Evolver system was supported by Graeme McCaig, Sara Salevati and Nathen Sorenson. As well from grants from NSERC Canada.

References

Ashmore, L., and Miller, J. (2004). Evolutionary art with cartesian genetic programming. Technical Online Report. http://www.emoware.org/evolutionary_art.asp.

Bentley, P., and Corne, D. eds. (2002). Creative evolutionary systems. Morgan Kaufmann.

Boden, M. (2003). The Creative Mind. London: Abacus.

Brown, J. (2009). Looking at Darwin: portraits and the making of an icon. Isis. 100(3), 542-70.

DiPaola, S., McCaig R., Carson K., Salevati S., & Sorenson N. (2013). Adaptation of an autonomous creative evolutionary system for real-world design application based on creative cognition. Proceedings of International Conference on Computational Creativity, 46-56.

DiPaola, S., and Gabora, L. (2009). Incorporating characteristics of human creativity into an evolutionary art algorithm. *Genet Prog and Evolvable Machines*, 10(2), 97–110.

Gabora, L. (2000). Toward a theory of creative inklings. In (R. Ascott, Ed.) Art, Technology, and Consciousness, Intellect Press, Bristol, UK.

Gabora, L. (2002). The beer can theory of creativity, in (Bentley, Corne, Eds.) Creative Evolutionary Systems, 147-161. San Francisco, CA.: Morgan Kaufmann.

Gabora, L., and DiPaola, S. (2012). How did humans become so creative? Proceedings of the International Conference on Computational Creativity, 203-210. May 31 - June 1, Dublin, Ireland.

Jennings, K. E. (2010). Developing creativity: Artificial barriers in artificial intelligence. Minds and Machines, 1-13.

Koza, J R., Keane, M A., and Streeter, M J. (2003). Evolving inventions. Scientific American. February. 288(2), 52 – 59.

Luo, J., and Knoblich, G. (2007). Studying insight with neuroscientific methods. Methods, 42,77-86.

Miller, J., (2011). Cartesian Genetic Programming, Springer.

Neisser, U. (1963). The multiplicity of thought. British Journal of Psychology, 54, 1-14.

Padian, K. (2008). Darwin's enduring legacy. Nature, 451(7179), 632–634.

Piaget, J. (1926). The language and thought of the child. Routledge and Kegan Paul, London.

Pease, A. and Colton, S. (2011). On impact and evaluation in computational creativity: a discussion of the turing test and an alternative proposal. Proc: AISB Symposium AI and Philosophy.

Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. Minds and Machines 17.

Sloman, S. (1996). The empirical case for two systems of reasoning. Psychol. Bull. 9(1), 3–22.

Yu, T., and Miller, J. (2001). Neutrality and the Evolvability of Boolean function landscape.