Chapter 10
CGP, Creativity and Art

Steve DiPaola & Nathan Sorenson

This chapter looks at evolved art and creativity Cartesian Genetic Programming (CGP). Besides an overview of evolutionary art, we discuss our work in modelling of artistic creativity based on the notion of contextual focus, which is the tendency for creative individuals to exhibit both focused concentration on a precise goal, as well as broad, associative though processes, which produce radical departures from convention. The model's implementation relies on the use of Cartesian Genetic Programming, which provides the essential property of genetic neutrality, which permits contextual fluidity. The model is used to generate creative portraits of Darwin, which serve to illustrate the focused and exploratory aspects of the creative process.

10.1 Introduction

A new field that has emerged over the last ten years in artificial intelligence systems is creative evolutionary systems. Creative evolutionary systems is used to evolve aesthetically pleasing structures in art, music and design. Within computer visual art, these systems are often referred to as evolutionary art systems.

The creative evolutionary systems research discussed in this chapter is based on Ashmore and Miller's work [21], which uses Cartesian Genetic Programming (CGP) first developed by Miller. CGP uses typical Genetic Programming (GP) evolutionary techniques (crossover, mutation, and survival), but has many features that we make document of in this chapter that allow the GP system to favour creative solutions over optimized solutions. Portrait painting was chosen for this project as it limits the creative space of all art paintings, weighs towards resemblance (similarity), and has a known portrait sitter/painter relationship well suited to explore computer creativity. Our system uses the same approach as Miller has discussed in Chapter 2, that is, the system generates portraits images using CGP where the two dimensional image pixel coordinates are the inputs to a CGP genotype that produce three 0 to 255 color outputs for each pixel in the two dimensional image array. In our case, those outputs are in HSV (Hue, Saturation and Value) space and our main contribution is in adding an autonomous creative fitness function to the system based on cognitive science research on human creativity.

According to Bentley from his seminal book on the subject [18], a creative evolutionary system is designed to 1) aid our own creative process and, 2) generate results to problems that traditionally required creative people to find the solutions. Bentley goes on to state that in achieving these goals, a creative evolutionary system may also appear to act 'creatively' - although this is still a source of debate. Unlike general evolutionary computation systems, creative evolutionary systems have been criticized because most of these systems use the presence of a human (often playing the role of the creative decision maker or fitness function) to guide the direction of the evolutionary search. Our CGP based portrait painter system specifically uses an automatic fitness function, albeit specific to a portrait painting where a portrait sitter resemblance is encouraged, thereby attempting to work through the human fitness function dilemma and directly explore how computer algorithms can be autonomously creative.

10.2 Creativity and Art

Certainly, creativity is a broad and complex notion that does not permit simple characterization. Creativity can be seen as a quality pertaining to both historical movements and solitary events, as a defining characteristic of both societies and individual, with understandable contention regarding the degree to which these instantiations are governed by the same phenomenon. Despite the inherent difficulties in constructing a comprehensive definition, much progress has been made in the characterization of creativity as an associative process. Indeed, Poincare's famous metaphor of disparate ideas that "rise in crowds" and "collide until pairs interlock" [2] seems only to be affirmed by recent work in neuroimaging technologies: an increase in associative brain activity can be seen during moments of creative thought, as new neurological connections between association cortices are formed [21]. This dynamic associative process is certainly linked with the ability that creative individuals have to make surprising and unanticipated departures from existing modes of thought.

In fact, case-study research demonstrates that creators often work in highly structured domains with well-specified rules against which they ultimately rebel [27, 18]. The traditional forms of the portrait, the sculpture, or the symphony become a point of departure for the truly creative artist. This is not to say Picasso, Michelangelo, and Stravinsky were not masters of the traditional forms; indeed it is precisely their mastery of existing techniques which made their innovations possible. In this sense, creativity is not simply the capacity to eschew traditional modes, but rather the ability to internalize and master them while still making associative connections that were previously not possible. This broad, associative imagination and narrow, focused mastery are two opposite impulses which characterize the creative thought process. As Feist states, "It is not unbridled psychoticism that is most strongly associated with creativity, but psychoticism tempered by high ego strength or ego control. Paradoxically, creative people appear to be simultaneously very labile and mutable and yet can be rather controlled and stable" [22]. Many theorists recognize the existence of comparable types of thought [12, 28, 14, 25]. Furthermore, many suggest that the ability to easily transition between these two modes is a defining characteristic of creativity [18, 29, 3]. This fluidity of mind is termed contextual focus [29], and requires both focused attention (typically linked with abstract thought and logical deduction) and a broad, expansive perspective (suited to the apprehension of unexpected correlations).

This dynamic contextual focus is the focus of this paper. In order to explore the nature of computational creativity, we develop a system to evolve artistic portrait paintings and implement a model of contextual focus to generate and evaluate artwork automatically. Our creative system simultaneously follows the precise and highly structured goal of representation, as well as the vague and associative notion of aesthetic quality.

10.3 Evolutionary Systems and Creativity

It is evident that the domain of artistic expression is particularly well suited to questions of machine creativity, as the standard techniques of domain-agnostic artificial-intelligence search do not apply. The process of creating art has no well-defined expected outcome; one cannot generate a creative artwork in the same way one can search for an effective chess move or compute an optimal load-bearing bridge design. There is simply no readily identifiable "problem" to be solved. The process used to formulate this problem definition is analogous to the artists struggle to realize an underlying vision to guide a work of art. This lack of a clear problem specification is exactly the sort of issue that both necessitates creative thought and makes its presence most evident.

Systems such as Harold Cohen's AARON [12] and Karl Sims genetic images [13] have popularized the notion that machines can autonomously produce output of aesthetic value. However, critics argue that the output is simply a function of the creativity of the system's designer, and not truly located within the machine. Indeed, how can the recognizable style of AARON be attributed to the machine and not its creator? Similarly, the evolution of the images produced by Sims' genetic system is directly guided by human interaction. Before we can claim to have embedded any degree of creativity within an automated system, it must be shown that the designers and users of that system are not ultimately responsible for the aesthetic decisions the system makes. We attempt to achieve autonomy by explicitly modelling the psychological process of contextual focus as a central component of our creative painter. It, therefore, exhibits the ability to evaluate its own designs in both a focused mode, with an adherence to a specific and well-defined goal, and also a broad, imaginative mode where more flexible judgment criteria are employed.

10.4 Evolutionary Art

Speaking broadly, creative evolutionary systems that combine with the aesthetic decisions of a human to judge fitness started well before computers. Standard historical selective breeding practises, where a human selects the parents for each generation from a given evolved set of choices, is the basis for centuries of 'creatively' modified trees, roses, corn, dogs, cats, cows and so on. Current evolutionary art systems borrow from this time tested approach. It was evolutionary biologist Richard Dawkins who first showed with his “Biomorphs” program that accompanied his 1986 book “The Blind Watchmaker” [14] that a computer can be combined with the aesthetic preference of a user to generate interesting results. Dawkins work inspired artists such as William Latham and Stephen Todd [15] as well as Karl Sims [13].

Karl Sims' work went on to inspire many of the modern evolutionary artists today. In his 2D work [13], Sims used a very rich instruction set, containing image processing functions as well as mathematical functions based on Lisp expression trees. As with most evolutionary art systems to follow, Sims system evolved a number of images (16 in his case) and allowed the viewers to pick their favourites, thereby allowing the most 'aesthetically pleasing' images to survive and mutate to the next generation. Other well known artists used similar techniques: Steve Rooke [16], also working in Lisp, is very well known for his artwork which added evolvable fractals to the function set; and Penousal Machado [17], a researcher at the Artificial Intelligence Laboratory at University of Coimbra, in contrast to Sims' complex function set, used a very simple function set which is believed to open up the possible search space.

These systems, as with most creative evolutionary systems, use a human (often the artist or viewer under interactive control) to make the aesthetic decisions after each evolutionary generation. In contrast, the use of an automatic fitness function which is able to make qualitative judgements constitutes a uniquely challenging research problem. Because of this, automatic fitness functions in evolutionary art are less the norm like they are in the general field of evolutionary computing However, a number of researchers are beginning to explore art based automatic fitness strategies: Bentley [18] in the design space; Miller and Thompson [19] in the field of electronic circuit construction; and John Koza, on his creative invention machine [20]. Automatic fitness functions for art, however, are especially difficult, and systems that use creative fitness functions in art are still quite naïve. Ashmore and Miller [21] have attempted to use an automatic fitness function with Cartesian GP that preferences images that contain circular objects (detected with a Hough Transform) or exhibit a high degree of complexity. However, this function only initializes a population and must defer to a human user’s input for further evolution. They also attempted to employ an automatic function for evolving towards a source image. We have based our system upon this notion of visual resemblance with a more sophisticated similarity function as well as adapting their system for a portrait painter process.

10.5 Genetic Programming and Creativity

GP's successes in producing novel and, arguably, creative designs are well publicized and implementing the creative process in an evolutionary algorithm such as CGP is conceptually pleasing, as the successive evolutionary stages of variation and selection map well to the engagement-reflection model of creative thought [22]. Engagement—the generation of possible ideas and solutions—manifests itself in the composition of simple building blocks defined by the GP's function set. Because these basic elements can be extremely basic and generic, the algorithm author need not inject too much pre-existing assumptions about anticipated solutions, which might restrict output to a certain type and therefore hinder creativity. The Reflection phase of creative practices is seen in the fitness function of the algorithm, and our model of contextual focus can be implemented as operations on the fitness function, modifying it to favour either precisely defined problems, or broad and vague notions.

The Classical GP approach poses, however, certain problems when tasked with modelling human creativity. Namely, GP excels at optimization problems that presume the existence of an optimal individual, which the search will then approximate. This is demonstrated by the fact that GP techniques typically use a single fitness function to evaluate every individual in every generation. The notion of an optimal individual is at odds with the process of contextually redefining and adjusting the goal of the search as it progresses. Indeed, problems that demand the use of creative intelligence do not have simple and stable evaluation criteria, and this is most certainly true of computer art, where the goal is not to produce an objectively optimal painting, but to explore variations and associations that are novel and unanticipated. Convergence to a particular individual solution halts exploration and stifles creativity. Indeed, regardless of the evaluation criteria used, if one individual in a population excels slightly better than the others, GP will tend to converge towards that value. By virtue of this process of optimization, diversity in the population is lost as all individuals assume the properties of the current leader. This loss of diversity has detrimental effects regarding the successful realization of the fluidity model of creativity: periods of narrow focus will damage the diversity of the population of solutions, essentially forgetting the individuals imagined during broad, associative phases. Returning to an associative phase from a narrow phase would essentially constitute starting from the scratch each time phases alternate. Clearly, this does not characterize the ease and fluidity that our model seeks to exhibit. Many solutions to the problem of maintaining diversity in evolutionary systems exist, and this is indeed a very well-researched subject. However, many solutions demand explicit organizational structures placed on them, such as sorting populations into different structural categories [23] or authoring a distance metric between individuals [24]. Such strategies rely on injecting a priori knowledge about the structure of the presumed solution into the system--something we wish to avoid.

10.5.1 CGP Advantages in Creative Systems

We use Cartesian Genetic Programming in our creative painting system as it is particularly well suited to avoiding these concerns. Though CGP shares with GP the same general process of iterative selection and variation, it differs in its representation of the genotype. The encoding is not a simple tree as in GP, but is rather a graph of indexed nodes. Each node represents a single basic function from the original function set, and can have a number of inputs and outputs. When this representation undergoes mutation, the connectivity of the graph is altered, possibly causing some nodes to become disconnected from the final program output. As certain nodes do not connect to the output, the information they represent becomes redundant. Genetic information becomes latent, and this gives rise to the essential property of CGP: neutrality. The very same solution, or phenotype, can be the result of a wide variety of different graphs, or genotypes. Thus, a narrow mode of evaluation can indeed focus temporarily on converging towards a single individual without necessarily invoking a permanent loss of genetic diversity, as a subsequent broad, associative phase of exploration will still have access to all the latent genetic information not visibly expressed in the phenotypes of the population. Furthermore, not only does CGP preserve diversity, it allows us to encourage such latent diversity explicitly. For example, in one set of experiments we implement the following rule: if the fittest individual of a population is identical to an individual in the previous generation for more than three iterations, the system chooses other genotypes that map to this same phenotype in favour of the current non-progressing genotype, thus promoting diversity in the latent genetic material.

10.6 Implementation

Our work is based on Ashmore and Miller's original application of CGP to genetic art [21]. Their basic approach, which we essentially follow and as outline in Chapter 2, consists of generating graphs with two inputs: the x and y coordinates of the pixel on the image plane, and three outputs: the hue, saturation and value (H,S,V) colour channels for that pixel. It should be noted that R,G, B outputs can be used and int eh case of Miller's initial work,either RGB or HSV could used. In our system, the HSV outputs better support artistic colour techniques in our contextual focus based fitness function in evaluating 2 of the 3 'rules of art' described in next section and figure 2 -- specifically dealing with tone (non colour graduation) separately which is represented in the V - value component and 2) evaluating warm and cool colour temperature ratios in the H – Hue component.

The functions in the function set (see Table 1) also can also use a random constant *param* as an input, which can be altered by mutation. The functions are kept intentionally simple and neutral to avoid imposing unnecessary structure into the ultimate results, allowing for a large search space. A graph of these functions constitutes an individual's genotype. When this compound function is evaluated for each pixel on an image plane, an image is produced, which is the individual's phenotype. It is this phenotype that is evaluated using our creative, contextually-focused, fitness function.

```
1: x | y;
2: param & x;
3: (x + y) % 255;
4: if (x>y) x - y; else y - x;
5: 255 - x;
6: abs (cos (x) * 255);
7: abs (tan (((x % 45) * pi)/180.0) * 255));
8: abs (tan (x) * 255) % 255);
9: sqrt ((x - param)2 + (y - param) 2); (thresholded at 255)
10: x % (param + 1) + (255 - param);
11: (x + y)/2;
12: if (x > y) 255 * ((y + 1)/(x + 1)); else 255 * ((x + 1)/(y + 1));
13: abs (sqrt (x - param2 + y - param2) % 255);
```

Table 1: Functions 1 through 5 use simple arithmetic operators on the x,y coordinates of the image. Functions 6 through 13 contain logical or trigonometric functions that are able to express more geometric shapes and colour graduations.

As with Ashmore and Miller’s work, the genotype is stored as an array of intergers of the length (n*4)+3 where n is the number of nodes, as seen in Figure 1. The last three integers in the chromosome are the output pointers for the Hue, Saturation and Value colour channels. The number of nodes in the chromosome affects the complexity of the output image. The greater the number of nodes will translate into more functions being used in defining the final image. In CGP each node normally defines the inputs to the node and the function only. Our functions are limited by having outputs between 0 and 255 for H, S and V. To increase the flexibility a further parameter has been added to each node which may or may not be used by the specified function.

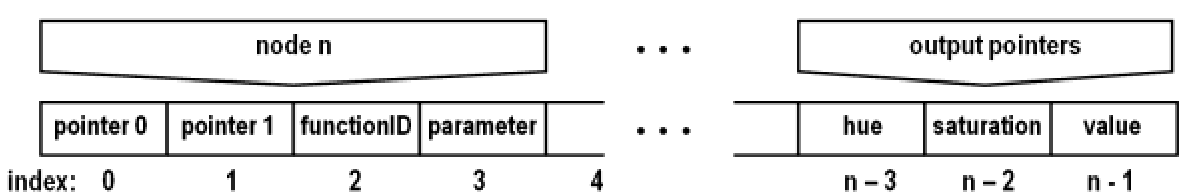

Fig. 1 Genotype schematic

Mutation occurs by calculating a random point along the chromosome, where a mutation rate specifies how many points are randomly chosen. Since along the chromosome some points can represent either a pointer, a function or a parameter, the mutation behavior has specific constraints for each of these three types.

Crossover points are selected between whole nodes so data within the nodes is retained. To create an offspring, the nodes before the crossover point come from one parent while the nodes after the crossover point, including the output pointers, will come from the other parent. A mechanism is in place that allows for cross over to occur when parents have genotypes of different lengths.

10.6.1 Fitness Function

Ashmore and Miller's original work consisted of initializing a population of interesting artworks, and allowing a human user guide subsequent evolution by evaluating the images. Our goal is to remove the human from the system by providing an artistic evaluation function that produces painterly portraits and exhibits contextual focus in its search for an aesthetic image. The goal of portrait painting is not to perfectly reproduce the appearance of the subject, (especially since the advent of photography), but rather to evoke a creative interpretation of the sitter. Therefore, the fitness function will, at times, emphasize the narrow and concrete goal of subject resemblance, while at other times defer to the fuzzy, associative and even contradictory “rules” of abstract art, with a psychologically-inspired model of contextual focus determining when to switch between them. CGP's phenotypic neutrality [30] ensures that the system does not destroy diversity when it seeks the narrowly defined goal of accurate resemblance; latent genetic material is still available for surprising associations later in the search process.

The fitness function determines resemblance by finding the mean-squared error between an image in the population and the source image. In our case, we take an image of Darwin as our subject. The abstract, painterly guidelines measure three different properties: the first is the composition of the face relative to the background; the second is the tonal similarity of the image as matched to a sophisticated artistic colour space model emphasizing warm-cool colour temperature relationships according to analogous and complementary colour harmony rules; and the third is the presence of a dominant and subdominant tone. These rules are drawn from the portrait painter knowledge domain, as detailed in [25].

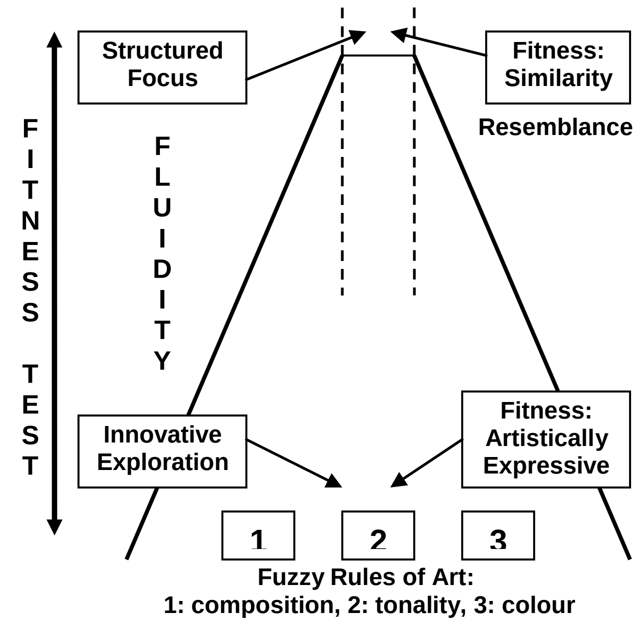

Fig.2 Our fitness function mimics human creativity by moving between restrained focus (resemblance) to more unstructured associative focus (resemblance + more ambiguous art rules of composition, tonality and colour theory).

The fitness function, then, calculates four scores for each image, (resemblance and the three painterly rules), using contextual-focus to inform the way in which these values are combined as the search progresses, as seen in Figure 2. Our model of contextual focus will alternate between emphasizing the highly structured goal of resemblance and encouraging the spontaneous exploration of the aesthetic principles of artistic composition.

10.6.2 Contextual Focus

By default, the fitness function favours resemblance by rating paintings using a ratio of 80% resemblance to a 20% non-linear combination of our three painterly rules. Several functional triggers can alter this ratio in different ways. For example, as long as a significant proportion of high-resemblance individuals exist in the population, in our case 80%, ("resemblance patriarchs"), the system will allow individuals with very high scoring under the painterly rules, ("strange uncles"), to be accepted into the next population. These individual with high painterly scores (weighted non-linearly to allow for those with a very high score in just one rule) are saved separately, and mated with the current population; if the system remains in this default state of focused resemblance, further offspring continue to be tested with the default 80% resemblance and 20% painterly rule test. Therefore, though we pull out and save these "strange uncles" to maintain artistic diversity, the focus of the genetic search is still towards resemblance.

The system, as a whole, will begin to favour the artistic rules when progress towards resemblance slows. As mentioned in Section 4.1, when a plateau, or local minima, is reached for a certain number of populations, the fitness function ratio beings to weight painterly rules higher than resemblance, on a sliding scale. Because artistic diversity has been explicitly encouraged in the focused resemblance phase, there is a great deal of artistically rich genetic information that is latent in the population. When the artistic rules are favoured over resemblance, this genetic information can manifest itself and a great deal of experimentation and exploration occurs.

Just as we saved artistically promising individuals during the focused stage, we are careful to isolate individuals with high resemblance during this artistic phase. These individuals are similarly allowed to pass onto the next generation when a certain proportion of aesthetically promising individuals is satisfied. Using this method, high resemblance individuals always remain in the population. When the resemblance of these individuals shows a marked improvement beyond the previous plateau, the system returns again to the default focused, resemblance mode.

10.7 Results

This system ran on one high-end PC for 50 days. Since the genes of each portrait can be saved, it is possible to re-combine (marry) and re-evolve any of the art works in new variants (Figure 3). As the fitness score increases, portraits look more like the sitter (Figure 4). This gives us a somewhat known spread from very primitive (abstract) all the way through realistic portraits. So in effect our system has two ongoing progressing processes: firstly, those portraits that pass on their resemblance strategies, making for more and more realistic portraits—the family ``resemblance patriarchs" (Figure 4), and secondly, the creative ``strange uncles," which are genetically related to the current resemblance patriarchs, but are exhibit greater aesthetic creativity. This dual evolving technique of patriarchs and strange uncles models the contextual focus of creative individuals as discussed in Section 1, that is the paradoxical technique where creative people use the existence of 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 (abstracted variation of the sitter image). That is, novel ideas require a pre-existing system to serve as a reference point from which innovation can occur.

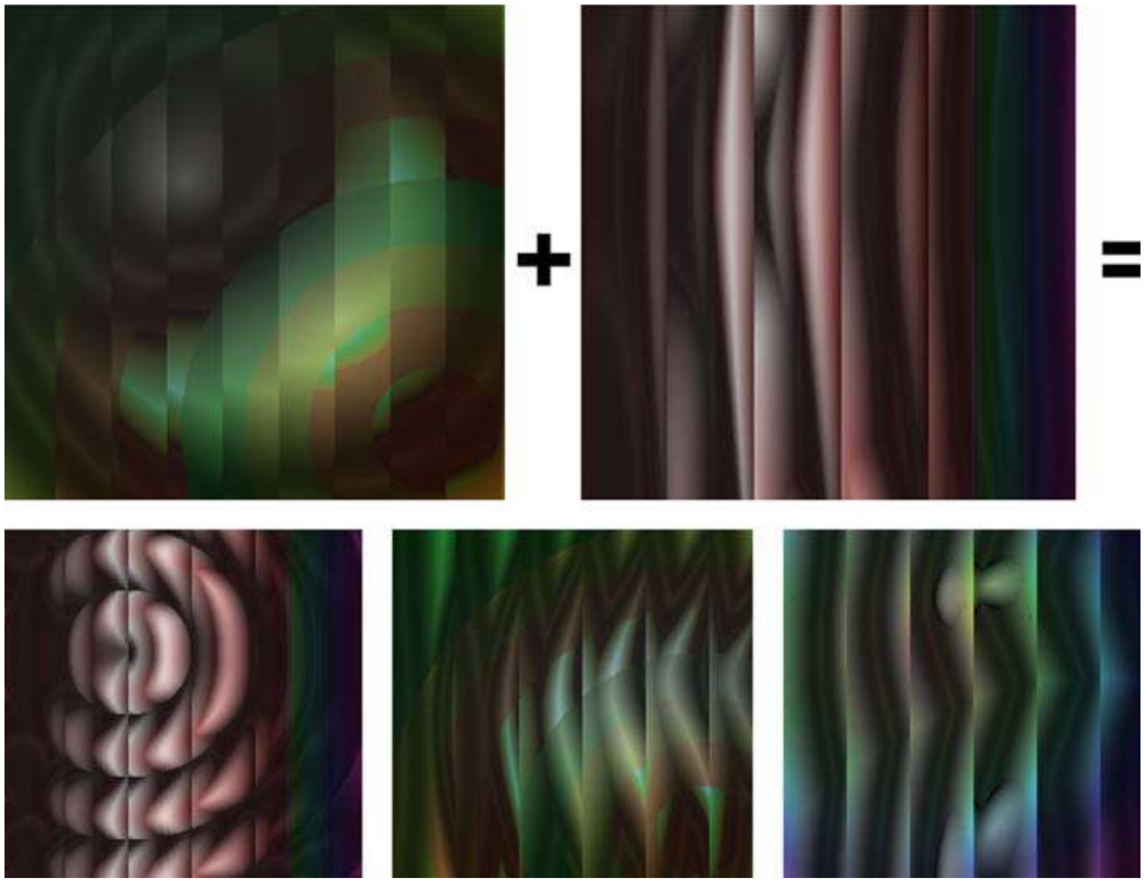

Fig. 3 Two portrait programs are mated together showing merged strategies of the offspring.

Fig. 4 Source Darwin portrait, part of the fitness function, followed by an evolved progression of portraits of best resemblance.

Another point worth highlighting is the difficulty of judging quality even in a population of equally poor paintings, which is currently an open research question regarding automated fitness functions [26]. Indeed, the initial population of images the system produces will bear absolutely no resemblance to the portrait subject, yet it is critical that the population is sorted in a precise and meaningful way, in order to guide evolution towards an aesthetic goal. To overcome this problem, we ensure that our fitness function effectively evaluates images in any stage of evolution. The associative contextual mode of our fitness function contributes greatly to achieving this generality, as, due to its broad nature, it is applicable to a likewise broad range of images. So, as opposed to direct resemblance, which only distinguishes between images meaningfully when there is a certain degree of resemblance already present, the abstract rules of colour and composition can be applied to virtually any image, regardless of how visually similar it is to its subject. Therefore, the system can switch its mental context at any point where it becomes difficult to distinguish between the images in its current population. Not only does this address the issue of ensuring effective evaluation at all stages of evolution, but it also models psychological creativity in a conceptually satisfying way: situations where there is no discernible way forward are precisely the times that call for creative exploration of alternatives.

However, it is ultimately those individuals that doggedly strive to resemble the Darwin image that move the system forward, as it is they who attain the highest resemblance scores and strategically move the system closer to the source image from in terms of resemblance. By allowing their related family members to be more innovatively artistic (via large local exploration) as safe variants from the patriarchs, we avoid the challenges to creativity that optimization presents, as discussed in Section 4. Figure 5 shows both types of individuals working synergistically, while Figure 4 only contains the resemblance patriarchs. We should emphasize that our goal is not to reproduce the Darwin portrait, but to explore a family tree of related and living portraits that inherit creative painting strategies through an evolutionary process. Ultimately, it is our hope to extend this system to be creative in a range of artistic and design oriented spaces beyond artistic portrait painting.

The images in Figure 5 show selected portraits in chronological order. These represent a larger collection, and show both those best at resemblance, as well as those that are artistically compelling. 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 seek, in different ways, to become better abstract portraitists.

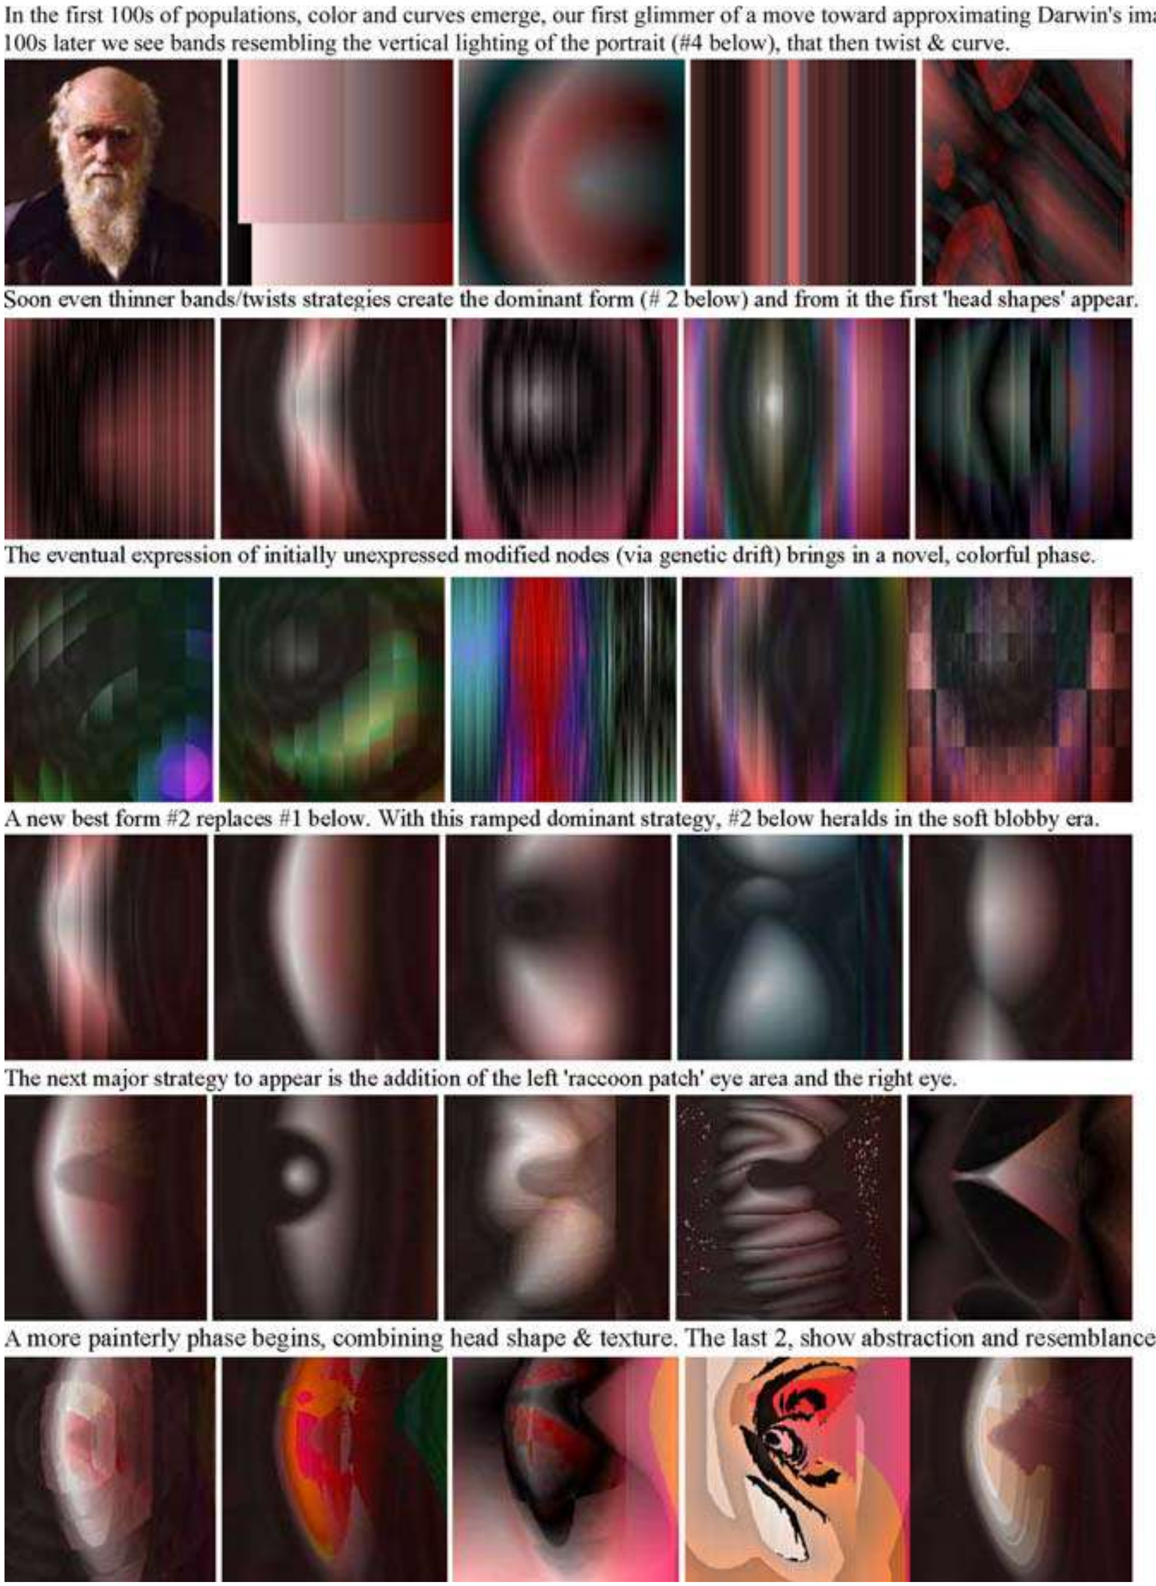

Fig. 5 Portraits in chronological order, selected as examples of the process (from a larger sampling at http://www.dipaola.org/evolve)

10.8 Conclusion and future direction

We have incorporated research on human creativity into the relativity new form of evolutionary computation, CGP, which has been successfully adapted to encourage the development of creative, painterly techniques. CGP exhibits genetic neutrality, which enables us to maintain the diversity needed to explore creative variations when faced with local minima. This technique proves to be well suited to the development of our contextual focus model of creativity, which requires the presence of such latent creative potential.

The domain of portrait painting was chosen because it leans heavily on resemblance (a closed and known issue for computer algorithms), but also has an opened ended creative element. As well, the portrait sitter to painter relationship is well suited to exploring computer creativity. The system indeed evolves creative strategies to become better abstract portraitists. We are continually refining the painterly portions of the automatic fitness function from lessons learned in past runs, and we are currently adding more creative, structural elements to this open-ended general system.

Key to this generality is increased understanding of how the potentiality of an idea changes and is affected by both the associative structure and the goals and desires of the mind it "finds itself in." To this end, future research will involve adding specific painterly and portrait knowledge with the goal of continuing to improve the automatic portrait painter system with human painterly knowledge. In addition, it is also possible (and possibly the direction of our next version of the program) to evolve the associative aesthetic fitness function simultaneously with the rest of the system. This can alter the dimensionality of the search space, the parameterization, as well as the representation of solutions, allowing for more creative automation.

Practically, to better approximate a human portraitist's technique, we are redesigning the functions in the function set to be reactions to the colour and position of the sitter image (the current system function set is blind to the sitter image, which is only used for evaluation). This way, any decision on a paint stroke output is a direct reaction to the input recognition (what the artist sees in the sitter scene). This would mean that, once a pleasing portrait image (individual) is created, the program could use its same painterly strategies on any new sitter image, thereby creating a true portrait painter.

Furthermore, a successful portraitist program might even have 'one-man' shows and take commissions, allowing its human creator to play a background role as its talent agent. It could eventually even be bred it with other successful portraitist programs similar to racing horses, allowing for experiments into cultural and collaborative creativity. This 'matching output stroke to input analysis' technique with other modifications would facilitate the realization of another goal: to have resolution-independent portraits, allowing small portrait sizes for speed during the evolving process, but larger sizes that reveal additional painterly and surface details for final artwork—as a human might make many creative sketches before the fully finished work.

We would like to explore the extent to which techniques used here can be transported to other domains such as art and design, music, authoring, HCI, entertainment, and gaming. The mechanisms will be kept general since we believe it is the associative, domain-general (rather than specialized, domain-specific) aspect of a creative architecture (organic or artificial) that is its greatest asset. Finally, we foresee a possible research application as a test bed for simulating creative processes or an educational tool for gaining hands-on understanding of evolutionary and creative processes.

10.9 Acknowledgements

We would like to thank Laurence Ashmore, Peter Bentley, Julian Miller, and James Walker for their correspondence, as well as Ashmore and Miller for their initial Java-based system that we adapted for our creative experiments. We would like to especially like to thank Liane Gabora, whose original ideas on creativity are a main part of this chapter.

References

1. H. Poincaré, Mathematical creation, in The Foundations of Science (Science Press, New York, 1913)
2. N. Andreasen, The Creating Brain: The Neuroscience of Genius (Dana Press, New York, 2005)
3. J. Feinstein, The Nature of Creative Development (Stanford University Press, Stanford, CA, 2006)
4. L. Gabora, The beer can theory of creativity, in Creative Evolutionary Systems, ed. by P. Bentley, D. Corne (Morgan Kaufmann, San Francisco, CA, 2002), pp. 147–161
5. G.J. Feist, The influence of personality on artistic and scientific creativity, in Handbook of Creativity, ed. by R.J. Sternberg (Cambridge University Press, Cambridge, UK, 1999)
6. T. Dartnell, Artificial intelligence and creativity: an introduction. Artif. Intell. Simul. Intell. Quart. 85 (1993)
7. U. Neisser, The multiplicity of thought. Br. J. Psychol. 54, 1–14 (1963)
8. L.J. Rips, Necessity and natural categories. Psychol. Bull. 127(6), 827–852 (2001). doi:10.1037/0033-2909.127.6.827
9. S. Sloman, The empirical case for two systems of reasoning. Psychol. Bull. 9(1), 3–22 (1996). doi: 10.1037/0033-2909.119.1.3
10. L. Gabora, Cognitive mechanisms underlying the creative process, in Proceedings of the Fourth International Conference on Creativity and Cognition, UK, 13–16 Oct. 2002, ed. by T. Hewett, T. Kavanagh, pp. 126–133
11. L. Gabora, Revenge of the neurds: characterizing creative thought in terms of the structure and dynamics of memory. Creat. Res. J. (in press)
12. H. Cohen, How to draw three people in a botanical garden in Proceedings of AAAI'88 (Morgan Kaufmann, San Francisco, CA, 1988) pp. 846-855
13. K. Sims, Artificial evolution for computer graphics. Comput. Graph. (ACM) 25, 319–328 (1991). doi: 10.1145/127719.122752

14. R. Dawkins, The Blind Watchmaker (W. W. Norton & Company Inc, New York, 1986)
15. S. Todd, W. Latham, Evolutionary Art and Computers (Academic, New York, 1994)
16. S. Rooke, Eons of genetically evolved algorithmic images, in Creative Evolutionary Systems, ed. By P.J. Bentley, D. Corne (Morgan Kaufmann, Los Altos, CA, 2002)
17. P. Machado, A. Cardoso, NEvAr - The Assessment of an Evolutionary Art Tool in G. Wiggins, (Ed.). Proceedings of the AISB00 Symposium on Creative & Cultural Aspects and Applications of AI & Cognitive Science, (UK, 2000)
18. P. Bentley, D. Corne (eds.), Creative Evolutionary Systems (Morgan Kaufmann, San Francisco, CA, 2002)
19. J. Miller, P. Thomson, Cartesian genetic programming, in Proceedings of the 3rd European Conference on Genetic Programming (Springer, LNCS 1802, 2000), pp. 121–132
20. J.R. Koza, M.A. Keane, M.J. Streeter, Evolving inventions. Sci. Am. 288(2), 52–59 (2003)
21. L. Ashmore, J. Miller, Evolutionary Art with Cartesian Genetic Programming. Technical Online Report, http://www.emoware.org/evolutionary_art.asp, 2004. Accessed 30 January 2005
22. M. Boden, The Creative Mind; Myths and Mechanisms (Routledge, New York, 2003)
23. M.A. Potter, K.A. De Jong, Cooperative coevolution: an architecture for evolving coadapted subcomponents. Evolutionary Computation 8(1), 1-29 (2000)
24. D.E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning, (Addison-Wesley, Boston MA, 1989)
25. S. DiPaola, Painterly rendered portraits from photographs using a knowledge-based approach, in Proceedings of Human Vision and Imaging Conference XII (International Society for Optical Engineering, San Jose, CA, 2007)
26. J. McCormack, Open problems in evolutionary music and art, ed. by F. Rothlauf et al., in Lecture Notes in Computer Science, vol. 3449, Proceedings of Applications of Evolutionary Computing, (EvoMUSART 2005), Lausanne, Switzerland, 30 March–1 April 2005 (Springer-Verlag, Berlin, Germany, 2005), pp. 428–436, ISSN: 0302-9743 & ISBN: 3-540-25396-3
27. E. Baker, Evolving Line Drawings, in Proceedings of the 5th International Conference on Genetic Algorithms, ed. by S. Forrest (Morgan Kaufmann, San Francisco, CA, 1993), p. 627
28. T. Dartnell, Artificial intelligence and creativity: an introduction. Artif. Intell. Simul. Intell. Quart. 85 (1993)
29 S. DiPaola, L. Gabora, Incorporating characteristics of human creativity into an evolutionary art algorithm, Genetic Programming and Evolvable Machines Journal, Vol 10, No 2, pp 97-110, (2009)
30. T. Yu, J. Miller, Neutrality and the Evolvability of Boolean function landscape, in Proceedings of the Fourth European Conference on Genetic Programming (Springer-Verlag, Berlin, 2001), pp. 204–217