Proceedings of the Eighth IASTED International Conference
COMPUTER GRAPHICS AND IMAGING
August 15-17, 2005, Honolulu, Hawaii, USA

A NEURO-GENETIC HYBRID MOTIF GENERATOR FOR GENETIC ART

James Wolfer
Indiana University South Bend
South Bend, IN, USA
email: jwolfer@iusb.edu

ABSTRACT

There have been a variety of methods for producing evolutionary art using Genetic Programming and other genetic algorithms. While some have included an underlying image, many of these systems produce aesthetically pleasing abstract images without overt structure. By using a physiologically inspired pulse-coupled neural network to find salient regions in an underlying image, and by subsequently introducing a motif function into the genetic programming system, we are able to augment the paradigm to introduce thematic image regions.

KEY WORDS

Computer Art, AI, Genetic Programming, Neural Networks

1 Introduction

Artists have used visual expression since recorded history. Ranging from early painting on cavern walls to computer animated motion pictures, humans use visual art to communicate information and emotion. Traditionally artistic expression has been through overtly intentional design with the artist mastering the mechanics of the currently available media and directing the conceptual expression within the physical constraints it imposed. This is true for the painter with the constraints of the paintbrush and canvas, and for the computer animator designing models for the screen.

Recent advances in science have lead to an interest in the psychology, sociology, and neural basis of artistic appreciation[1, 2]. Understanding the underlying neurological mechanisms will, hopefully, lead to a broader understanding of art and aesthetics. One aspect of these studies is considering, and in this case applying, insights from low-level neural investigations. Specifically, creating physiologically inspired saliency maps to introduce a visual motif into genetic art through a pulse-coupled neural network.

Finally, recent work in evolutionary computing has allowed artists to engage in a more passive form of artistic expression, allowing the artist gage the aesthetic merit of automatically evolved artforms. Examples include the work of Sims[3], Jones and Agah[4], Rooke[5], and Wiens and Ross[6]. More recently, studies have introduced methodology allowing the evolutionary system to self-assess the aesthetic appeal of the resulting images[7, 8].

While these systems often feature abstract images and textures, they have also included evolutionary seeding from input images[3, 8]. In this work we augment this evolutionary framework to selectively incorporate image features in a motif function for a genetic programming system. We implement this system in three basic steps:

1. Create a "salience image" from the output of a pulse-coupled neural network.
2. Based on the resulting salience image, define one or more functions to be made available to an interactive genetic programming system.
3. Execute the resulting system, using observer feedback as a fitness function.

2 Salient Image Regions

The detection of salient region and border pixels is an important step in many computer graphics, vision, and image processing applications including those in evolutionary computing[9]. Many approaches to isolating these regions have been described. Often they include low level image processing such as noise reduction followed by edge detection operators such as the popular Soble and Canny edge detectors[10]. Post-processing is then applied to extract model information from the edges and regions identified during the low level processing. Typically this post-processing incorporates a priori information about the expected forms in the image. Post-processing models include active shapes and contours, mathematical morphology, fractal analysis, edge and shape, and temporal models[11, 10].

3 The PCNN

Recently there has been some interest in more biologically inspired models in applied computer vision. One such model is the pulse-coupled neural network. The Pulse-Coupled Neural Network (PCNN) attempts to model neuron interactions in time. Based upon Eckhorn's work modeling neural interaction on the visual cortex of the cat[12, 13] and, more recently, primates, the PCNN forms a high order network in which spikes through time form a succession of binary outputs. When such a model is stimulated by the values of an input image the network produces a sequence of binary images reflecting the propagation of

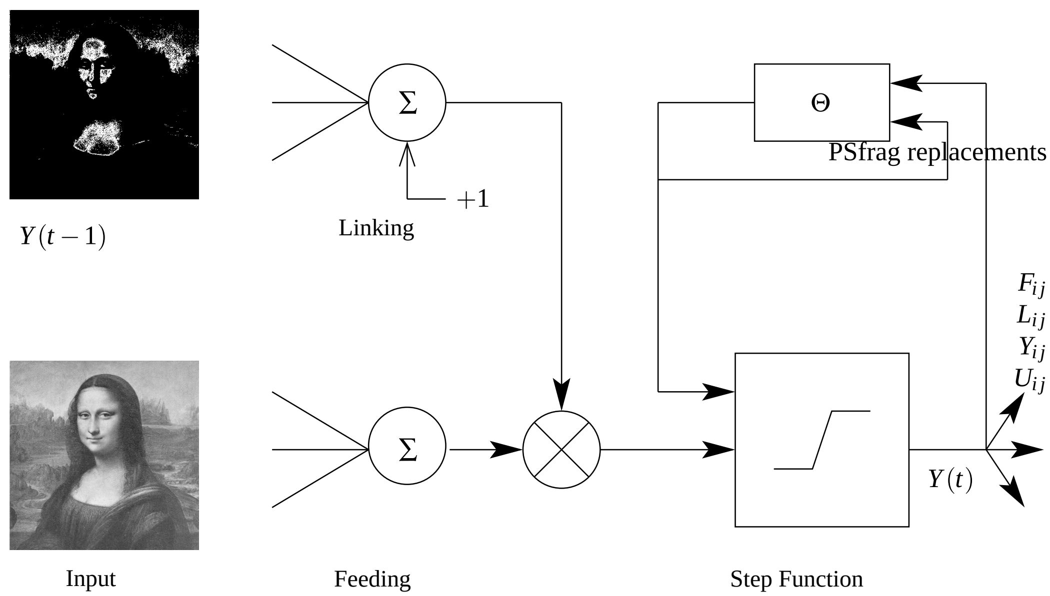

Figure 1. PCNN Schematic

local activity. Attractive aspects of the PCNN for genetic art include relative immunity to translation, scale, and rotation in the image[14]. Figure 1 is a schematic diagram of a single PCNN neuron. It is divided into three primary functions: feeding, linking, and pulse generation. This PCNN neuron is modeled as follows[15, 16]:

$$F_{ij}(t) = e^{-\alpha_F \delta t} F_{ij}(t-1) + S_{ij} + V_F \sum_{kl} W_{ijkl} Y_{kl}(t-1)$$

$$L_{ij}(t) = e^{-\alpha_L \delta t} L_{ij}(t-1) + V_L \sum_{kl} M_{ijkl} Y_{kl}(t-1)$$

$$U_{ij}(t) = F_{ij}(t)(1 + \beta L_{ij}(t))$$

$$Y_{ij}(t) = \begin{cases} 1 & \text{if} \quad U_{ij}(t) > \Theta_{ij}(t) \\ 0 & \text{Otherwise} \end{cases}$$

$$\Theta_{ij}(t) = e^{-\alpha_\Theta \Delta t} \Theta_{ij}(t-1) + V_\Theta Y_{ij}(t)$$

Where F is the feeding component, L the linking component, U the neuron internal activity, Y the neuron output, and Θ the dynamic threshold. M and W encode weights from the individual inputs along with a receptive field for the feeding and linking functions respectively, and β is the linking strength. These equations are applied in sequence at each iteration of the simulation.

In image processing, an individual neuron receives input to it's feeding function from a single, scaled, gray level pixel in the original image along with a receptive field consisting of a weighted neighborhood. This results in one artificial neuron being directly stimulated by a corresponding pixel and its neighbors from the input image, preserving geometric structure of the image.

Assuming that the threshold is initially set to zero, any activity at the input will cause a corresponding output from the pulse generation. This, in turn, raises the threshold suppressing subsequent output. As the threshold decays neurons with activity exceeding the threshold pulse, reestablishing a high threshold for them, but also raising the probability that adjacent neurons will be fire at the next iteration as a result of the linking feedback to the receptive field. In this sense each artificial neuron can be seen as initiating an autowave of activity which propagates until colliding with another wavefront.

Figure 2 illustrates the action of the PCNN when stimulated by the image of the Mona Lisa as shown in Figure 1. For this, and all subsequent images, alpha was empirically fixed at 10.0, 1.0, and 15.0 for the feeding, linking, and threshold computations respectively, and beta was 0.7.

Initially, at iteration one, every non-zero pixel causes the PCNN to fire driving the threshold high. Over time the threshold decays until neurons connected to those pixels providing the highest stimulation fire, in this case at iteration seventeen. This, in turn, stimulates the surrounding pixels causing them to fire if they are close to their respective thresholds.

The composite salience map is created by presenting an image to the inputs of the PCNN and iterating the PCNN. Each time a unit fires, after the initial spike at iteration 1, the time that it fired (i.e. the iteration number) is recorded in an array at the corresponding pixel location. The network continues iterating until there have been ten time steps in which one or more neurons fired. The resulting map is an image with increasing pixel values as the wave of activity propagates.

Figure 3 displays a contrast enhanced salience image consisting of ten firing times. The number of firings was chosen empirically after examining the PCNN wave propagation.

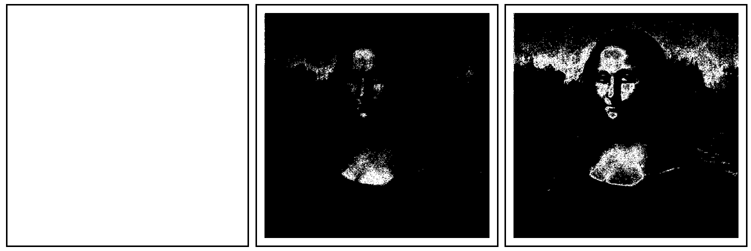

(a) Iteration 01 (b) Iteration 17 (c) Iteration 18

(d) Iteration 19 (e) Iteration 20 (f) Iteration 21

Figure 2. PCNN Iterations Mona Lisa.Note that all pixels initially fire in iteration one.

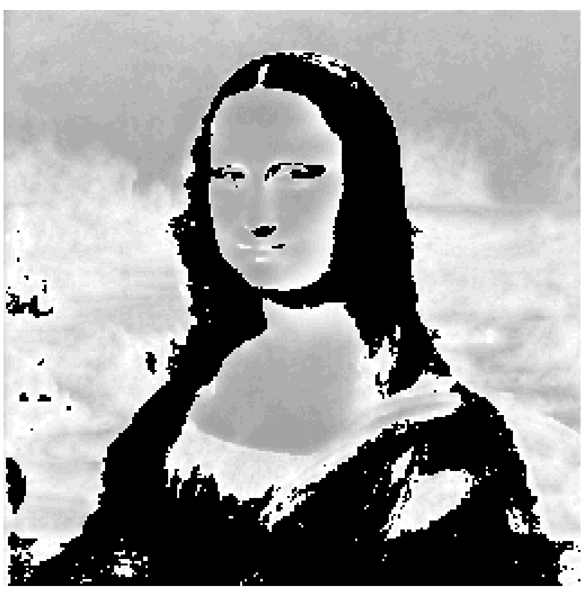

Figure 3. Salience Map

4 Artistic Applications

To incorporate the salience maps into our images we introduce a function into a Genetic Programming system patterned after that of Sims[3]. Since this is such a popular approach to genetic art we will only summarize it here. Fundamentally, this system generates a population of functions with parameters that may include the current location being evaluated. The functions are executed for each desired pixel in the output image, and the returned values are assembled into images and displayed. The fitness of the functions are then evaluated, in this case interactively by a human operator based on the perceived aesthetic properties of the resulting images, and those selected become raw material for a new generation of functions.

While there are a variety of ways to apply the salience map, we provided a new function, which we call motif. Motif is a function of the composite saliency map at the current location and one data value passed by the current Genetic Program individual. While several possibilities were explored, the following resulted in a good compromise between structure and diversity:

$$\mathrm{Motif}(x,y,d) = \begin{cases} d^2 & \text{if} \quad d \leq SM(x,y) \\ d & \text{Otherwise} \end{cases}$$

Where $SM(x,y)$ is the salience map value at location x,y, and d is provided by the currently running genetic program. Additional primitive operations available to the genetic programming system include basic arithmetic operators, sin, cos, log, exp, square, sqrt, if, and a random number generator.

To illustrate the result we selected the four images shown in Figure 4. Figure 4a is the well known Mona Lisa, Figure 4b is a photograph of the Golden Record included on the Voyager spacecraft[1], Figure 4c is Mount Rushmore[2].

[1]Courtesy NASA/JPL-Caltech

[2]Courtesy U.S. National Park Service

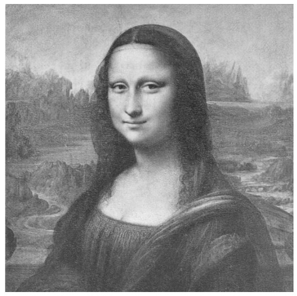

(a) Mona Lisa

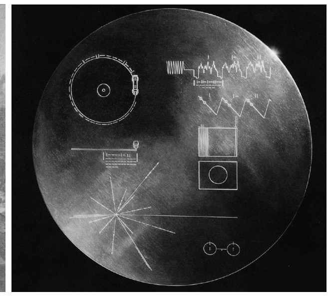

(b) Voyager Record

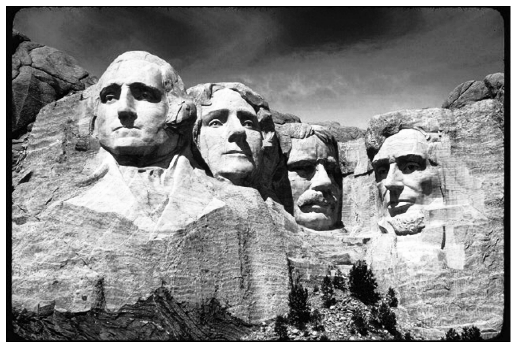

(c) Mt. Rushmore

Figure 4. Sample Images

With the exception of the Mount Rushmore image, which was contrast enhanced, these images have not been preprocessed.

Figure 5 demonstrates a variety of images evolved with the Mona Lisa as a motif and Figure 6 shows the blending of a random element and the Mona Lisa motif. We call the images in Figure 6 “petroglyph” since each is reminiscent of an ancient petroglyph. In addition to the motif imagery, the system can and does generate purely abstract images as illustrated in Figure 7.

The results from the Mount Rushmore image are shown in Figure 8 and range from nearly photographic to an imprint in random noise. Finally, the Voyager Golden Record is shown in Figure 9.

5 Observations and Future Work

While we chose to fix the PCNN parameters to illustrate its effectiveness over a variety of photographic images, optimal parameter selection is an open question and bears further study. The PCNN is also sensitive to noise and the results from the Mona Lisa in particular could be improved by preprocessing the image through a lowpass filter. This was not done in the interest of showing its applicability to a range of imagery. The only preprocessing was to contrast enhance the Mount Rushmore image since the original was a very low contrast image.

Since the primary objective for this work was structural we focused on monochromatic images. Color can, however, be introduced by a variety of methods including considering each band independently or processing CIE luminance images formed from the three primary color bands.

While the application presented here is interesting, it is only an example of one potential application of the resulting salience maps. It is likely that other applications, such as the swarm based, non-photorealistic rendering described by Semet et.al[9] could find this as an alternate salience detector. We also speculate that there may be enough residual information in these PCNN detected maps to estimate the universal metric based aesthetic fitness as described by Svangard and Nordin. We look forward to investigating these possibilities in the future.

6 Conclusion

We have applied the pulse-coupled neural network as a physiologically inspired salience detector, and have illustrated it use as a motif generator for genetic art. We believe that it has other application potential and hope to continue exploring that potential in the future.

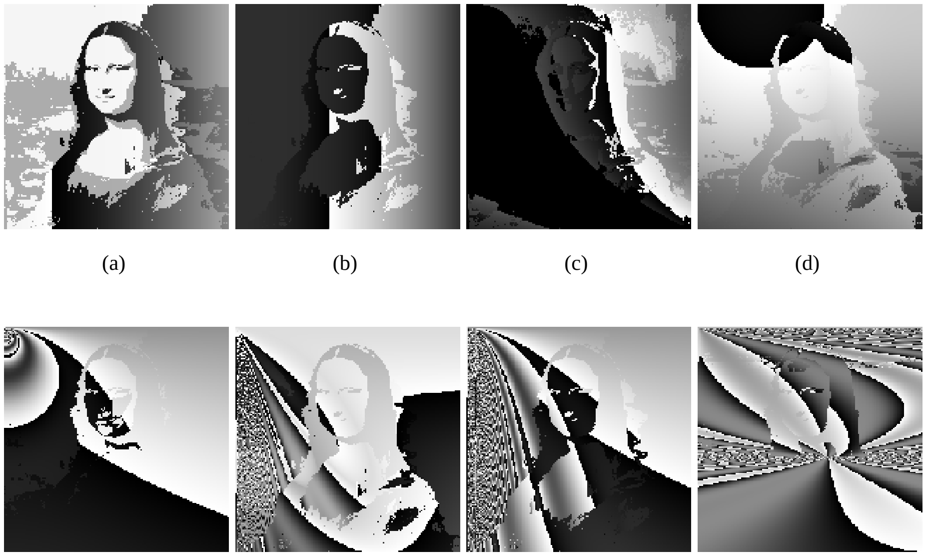

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 5. Mona Lisa

(a) (b) (c)

Figure 6. Mona Lisa “Petroglyph”

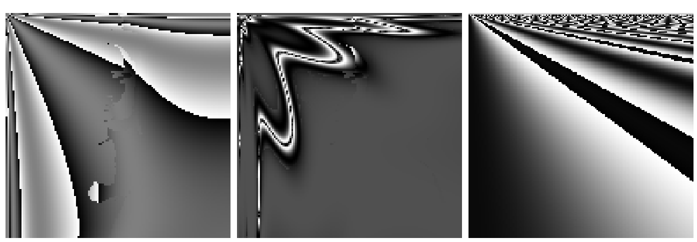

(a) (b) (c)

Figure 7. Abstract Results

References

[1] Semir Zeki. *Inner Vision: An Exploration of Art and the Brain*. Oxford Press, 2002.

[2] JoAnn Wypijewski. *Painting by Numbers: Komar and Melamid’s Scientific Guide to Art*. University of California Press, 1997.

[3] Karl Sims. Artificial evolution for computer graphics. In *Computer Graphics*, pages 319–328, 1991.

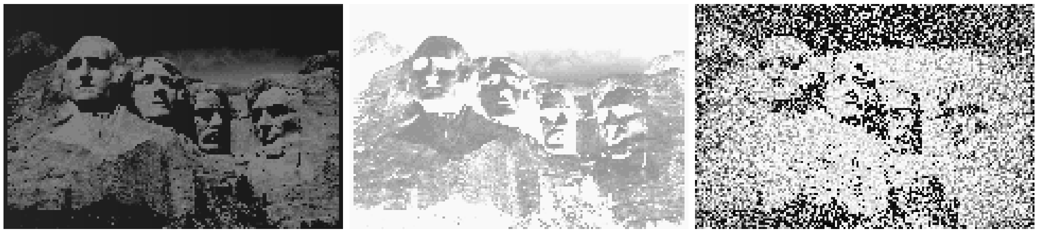

(a) (b) (c)

Figure 8. Mount Rushmore

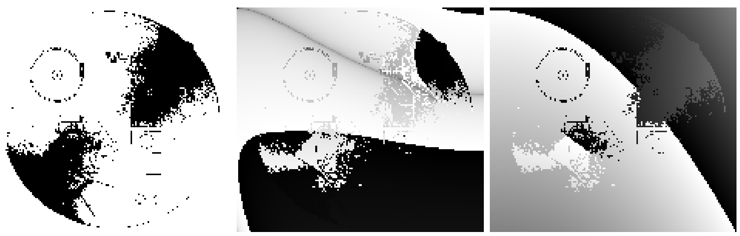

(a) (b) (c)

Figure 9. NASA Golden Record

[4] Mara Elizabeth Jones and Arvin Agah. Evolution of digital images. *IEEE Transactions on Systems, Man, and Cybernetics C*, 32(3):261–271, August 2002.

[5] Steven Rooke. *Eons of Genetically Evolved Algorithmic Images in Creative Evolutionary Systems*. Morgan Kaufmann, 2002.

[6] A. Wiens and B.J. Ross. Gentropy: Evolutionary 2d texture generation. *Computers and Graphics Journal*, 26(1):75–88, 2002.

[7] Nils Svangard and Peter Nordin. Automated aesthetic selection of evolutionary art by distance based classification of genomes and phenomes using the universal similarity metric. In *Applications of Evolutionary Computing, EvoWorkshops2004: EvoMUSART*, pages 447–456. Springer Verlag, 5-7 April 2004.

[8] P. Machado and A. Cardoso. Nevar – the assessment of an evolutionary art tool. In *AISB'00 Symoosium on Artificial Intelligence and Creativity in the Arts and Science*, 2000.

[9] Yann Semet, Una-May O'Reilly, and Fredo Durand. An interactive artificial ant approach to non-photorealistic rendering. In *Genetic and Evolutionary Computation - GECCO 2004*, pages 188–200, 2004.

[10] E.R. Davies. *Machine Vision: Theory, Algorithms, Practicalities*. Academic Press, 1997.

[11] Marian M. Choy and Jesse S. Jin. Improving border identification in two-dimensional echocardiograms using temporal information. In *IEEE Engineering in Medicine and Biology*, pages 879–880, 1996.

[12] Reinhard Eckhorn. Neural mechanisms of scene segmentation: Recordings from the visual cortex suggest basic circuits for linking field models. *IEEE Transactions on Neural Networks*, 10(3):464–479, May 1999.

[13] Reinhard Eckhorn, Alexander M. Gail, Andread Bruns, Andreas Gabriel, Basim Al-Shaikhli, and Mirko Saam. Different types of signal coupling in the visual cortex related to neural mechanisms of associative processing and perception. *IEEE Transactions on Neural Networks*, 15(5):1039–1052, September 2004.

[14] John L. Johnson. Pulse-coupled neural nets: translation, rotation, scale, distortion, and intensity signal invariance for images. *Applied Optics*, 33(26):6239–6253, September 1994.

[15] Thomas Lindblad and Jason M. Kinser. *Image Processing using Pulse-Coupled Neural Networks*. Springer Verlag, 1998.

[16] John L. Johnson and Mary Lou Padgett. Pcnn models and applications. *IEEE Transactions on Neural Networks*, 10(3):480–496, May 1999.