Automated Evolution of Interesting Images

Joshua E. Auerbach[1]

[1]Morphology, Evolution & Cognition Laboratory
Department of Computer Science, University of Vermont
Burlington, VT 05405
joshua.auerbach@uvm.edu

Extended Abstract

Recent work (Secretan et al., 2011) has demonstrated that it is possible to evolve interesting images produced by Compositional Pattern Producing Networks (CPPNs) (Stanley, 2007) through interactive evolution. However, interactive evolution is a slow process that requires the active involvement of human users. It is desirable to evolve interesting images without requiring human users to perform selection. In this work I explore alternate methods of evolving interesting images from CPPNs that are completely automated, yet in some cases still indirectly informed by what humans find interesting.

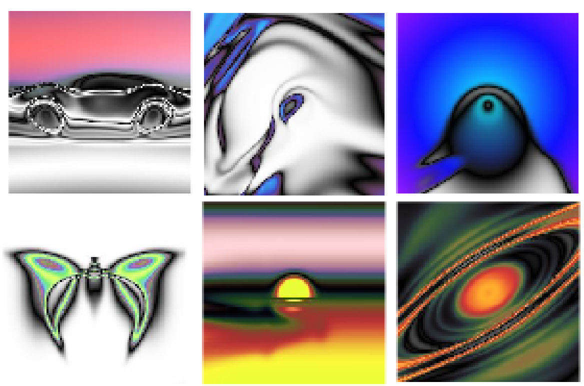

Figure 1: Sampling of interesting and familiar looking images produced by Picbreeder. Taken from (Secretan et al., 2011).

Picbreeder (`http://www.picbreeder.org`) is a website created by researches in the EPLEX group at the University of Central Florida for the purpose of interactively evolving images. This system has produced many interesting and familiar looking images. A small sampling of these can be seen in Figure 1. The images in Picbreeder are produced by CPPNs, a form of indirect encoding that abstracts biological development to produce outputs with the familiar biological properties of symmetry, repetition and repetition with variation. Commonly CPPNs are evolved via CPPN-NEAT, an extension of the state of the art NeuroEvolution of Augmenting Topologies (NEAT) (Stanley and Mikkulainen, 2001) algorithm, applied to CPPNs. Picbreeder employs this algorithm with selection based on the interactive choices of human users over the internet. Likewise, this algorithm (or variants of it) are employed in the work presented here[1].

In lieu of employing human users in the experiments conducted in this work, the fitness of individual images is calculated automatically and selection is based on these calculated fitness values. 400×400 pixel images are evolved. The two primary ways in which these images' fitnesses are evaluated are their complexity, defined in terms of the size of a zlib compressed representation of the image and their ability to maximize the number of results returned when the image is used as a search query to Google's search by image (SBI) feature[2].

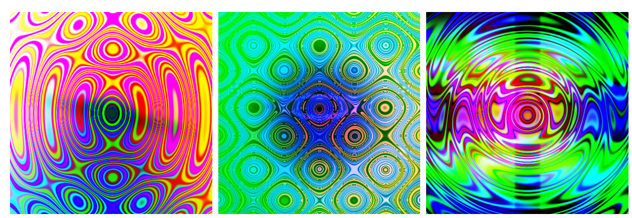

Figure 2: Sampling of images evolved to be maximally complex. These images evoke 1960s concert posters.

The initial hope was that by maximizing the number of results when querying SBI it would be possible to indirectly capture notions of interestingness, because images which are interesting to humans are precisely those likely to exist on the internet and be indexed by Google. However, when the sole objective of evolution was to maximize this number evolution tended to converge on images that were entirely one color (though the specific color varied across evolutionary trials). It is hypothesized that this is

[1]The selection mechanism is what differentiates the current work from Picbreeder. The CPPNs evolved in this work make use of the same inputs and outputs as those in Picbreeder including the use of the Hue Saturation Brightness (HSB) color space.

[2]`http://www.google.com/imghp?sbi=1`

because searching on a single color will find many images containing that color, but the details of SBI are proprietary and unknown to the author.

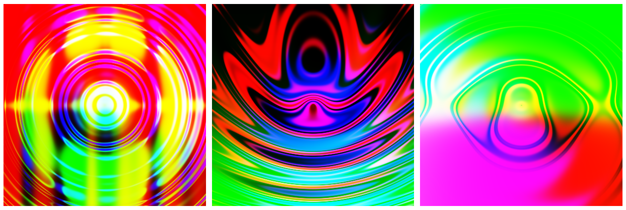

Figure 3: Sampling of images evolved to maximize complexity and SBI hits combined multiplicatively into a single fitness function.

Alternatively, complex images tend to be interesting so by selecting for maximally complex images it should be possible to produce interesting results. A sampling of images selected to be maximally complex is shown in Figure 2. Finally, by combining these objectives either multiplicatively or through pareto based multi-objective selection[3] it should be possible to evolve images both complex and informed by the corpus of images that exist on the web. Images evolved under both these schemes are shown in Figures 3 and 4 respectively.

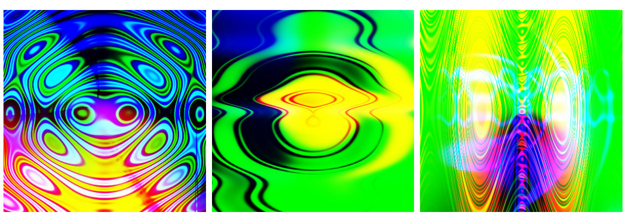

Figure 4: Sampling of images evolved to maximize complexity and SBI hits using pareto based multi-objective optimization.

While none of these selection mechanisms produced images with the familiarity of those produced by Picbreeder the results are still interesting to look at for artistic purposes. Modifying the evolutionary search in other ways, such as by allowing recurrent connections within the CPPN genomes may allow for the creation of images that are interesting in distinct ways and is beginning to be investigated (see Figure 5 for examples). Additionally it is possible that by using a larger number of evaluations[4], by combining these fitness criteria with others, or by incorporating keywords into the SBI searches it may be possible to automatically evolve images that are both interesting and familiar without resorting to the direct user evaluations employed in Picbreeder.

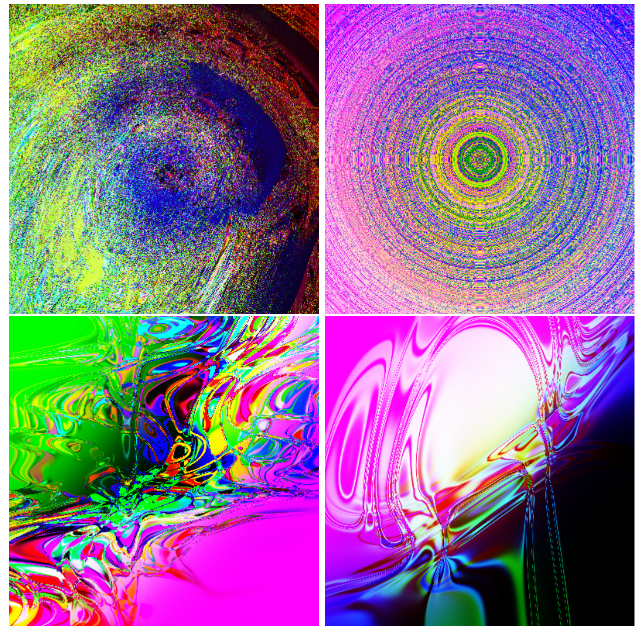

Figure 5: Sampling of images evolved using recurrent CPPN connections.

Acknowledgements

This work was supported by National Science Foundation Grant PECASE-0953837.

The author would like to thank Peter Andrews, Jeff Clune, and Jean-Baptiste Mouret for a conversation that led to the development of the ideas in this paper.

The author also acknowledges the Vermont Advanced Computing Core which is supported by NASA (NNX 06AC88G), at the University of Vermont for providing High Performance Computing resources that have contributed to the research results reported within this paper.

References

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6:182–197.

Secretan, J., Beato, N., D'Ambrosio, D. B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J. T., and Stanley, K. O. (2011). Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation Journal*, pages 373–403.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2):131–162.

Stanley, K. O. and Miikkulainen, R. (2001). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:2002.

[3] For this purpose CPPN-NEAT was modified to use a selection mechanism based on NSGA-II (Deb et al., 2002)

[4] Here the number of evaluations was rate limited because querying Google too frequently will cause one's IP address to be blocked.