Adversarial evolution and deep learning – how does an artist play with our visual system?

Alan Blair

School of Computer Science and Engineering
University of New South Wales, Sydney, Australia
blair@cse.unsw.edu.au

Abstract. We create artworks using adversarial coevolution between a genetic program (HERCL) generator and a deep convolutional neural network (LeNet) critic. The resulting artificially intelligent artist, whimsically named Hercule LeNet, aims to produce images of low algorithmic complexity which nevertheless resemble a set of real photographs well enough to fool an adversarially trained deep learning critic modeled on the human visual system. Although it is not exposed to any pre-existing art, or asked to mimic the style of any human artist, nevertheless it discovers for itself many of the stylistic features associated with influential art movements of the 19th and 20th Century. A detailed analysis of its work can help us to better understand the way an artist plays with the human visual system to produce aesthetically appealing images.

Keywords: evolutionary art, AI-generated art, artist-critic coevolution, adversarial training, computational creativity

1 Introduction

There has recently been renewed interest in the paradigm of artist-critic coevolution or adversarial training in which an artist (generator) tries to produce images which are similar in some way to a set of real images, and a critic tries to discriminate between the real images and those generated by the artist [1, 2].

The earliest work in this area followed an interactive evolution scenario, with a human playing the role of the critic, and the artist trained by some form of evolutionary computation such as Biomorphs [3], Genetic Programming [4], Cellular Automata [5] or Compositional Pattern Producing Networks [6]. In these systems, several candidate images appear on the screen and the user is invited to select one or more of them for inclusion in the next generation. These approaches have produced some remarkable images, but the process can be time-consuming for the human as several dozen generations are often required in order to produce a pleasing image.

To save human effort, the critic has sometimes been replaced with a pre-trained image or face detection system [7, 8], or a Convolutional Neural Network

(CNN) trained to mimic human preferences [9]. Similar machine learning approaches have been used to explore aesthetic feature selection [10, 11] or to model collective artistic behavior using Self Organizing Maps [12].

An exciting new paradigm was introduced in 2008 where images are created through an adversarial arms race between a creator (generator) and an adaptive critic [13]. The critic is rewarded for its ability to distinguish "real" images from those generated by the creator, while the creator is rewarded for generating images that will fool the critic into thinking they are real. Typically, a genetic algorithm was used for the creator, while the critic was a 2-layer neural network trained to classify images based on certain statistical features extracted from the image [14–16].

A powerful new variant of this adversarial paradigm was developed in 2014 known as Generative Adversarial Networks (GANs) [17], where the generator and discriminator (critic) are both CNNs trained by gradient descent. Compared to other approaches, GANs produce astonishingly realistic images [18], and also have the advantage that a single trained network can produce a great variety of images from latent variables.

The aim of the present work is not to generate photo-realistic images, nor to copy existing artistic styles, but rather to create new abstract artworks sui generis. Our system employs a recently introduced hybrid approach [19] where the generator is a Genetic Program evolved by hierarchical evolutionary recombination (HERCL) [20], while the critic is a GAN-style convolutional neural network (LeNet) [21] trained by gradient descent. Each image is generated via a genetic program which takes as input the x and y coordinates of a pixel, and produces as output the R,G,B values assigned to that pixel in the image [22].

Because it is comprised of a HERCL generator and a LeNet critic, the resulting adversarial art generation system is sometimes whimsically referred to as if it were a real artist named "Hercule LeNet".

2 Interplay between Generator and Critic

It is important to note that this hybrid evolution and deep learning system is not trained on any pre-existing artworks, nor is it asked to mimic the style of any human artist, nor to directly satisfy the whims of any human observer. Rather, it is presented with a number of real images, and aims to produce synthetic images which have low algorithmic complexity (due to selective pressure for shorter programs), but which nevertheless resemble the real images well enough to fool an adversarially trained (LeNet) critic with an architecture modeled on the human visual system.

This very broad objective effectively liberates the artificial artist from stylistic preconceptions, allowing it to freely explore and develop its own artistic style(s), which can then be analysed.

In the context of image classification, it has been pointed out in a number of contexts that deep convolutional networks are easily fooled [8]. This is certainly a problem if accurate recognition is the goal, or if the networks are fooled in ways

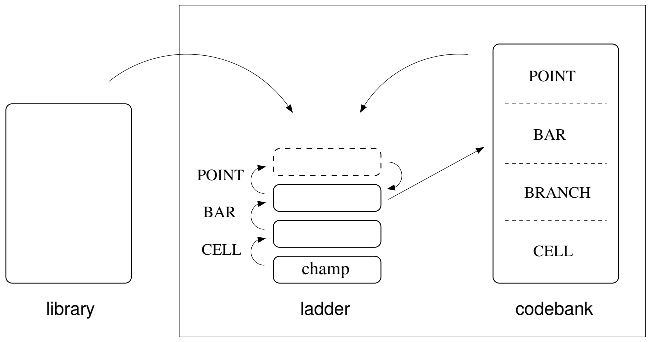

Fig. 1: Hierarchical evolutionary re-combination. If the top agent on the ladder becomes fitter than the one below it, the top agent will move down to replace the lower agent (which is transferred to the codebank). If the top agent exceeds its maximum number of allowable offspring without ever becoming fitter than the one below it, the top agent is removed from the ladder (and transferred to the codebank). Code from related tasks (in this case, the code that generated earlier images in the gallery) is kept in a library and made available for genetic recombination.

which appear puzzling to a human observer. However, if the networks are fooled in ways which seem quite familiar to us as humans, this might indicate that the networks share certain quirks of the human visual system – quirks which can be exploited in interesting ways by abstract visual artists.

When it comes to impressionist, minimalist or abstract art, we as humans are not actually deceived into thinking that the images are real; rather, our appreciation of these artworks relies at least in part on their ability to challenge and tantalize our visual system, prompting us to look at things from a new perspective. If our evolutionary artist (HERCL) is able to fool the artificial visual system of the convolutional network (LeNet) in an analogous manner, it may produce images which have a similar appeal for human observers, and help us to gain a better understanding of visual aesthetics.

3 Adversarial Training Paradigm

Our adversarial training paradigm, essentially the same as in [19], is as follows: In the first round, the LeNet critic is trained to assign a low cost (close to 0) to all the real images and a high cost (close to 1) to a blank (i.e. completely white) image. Training is by stochastic gradient descent using the Adam optimizer [23]. In all subsequent rounds, a HERCL artist is evolved with the aim of producing a new image to which the current critic will assign as low a cost as possible — continuing until the cost becomes lower than 0.01, or a maximum of 200000 images have been generated. The resulting minimal-cost image is then added to a "gallery" of images, and a new critic is trained to assign a low cost to the real images and a high cost to all the images in the gallery (i.e. the minimal-cost image produced by the artist in every previous round). This process continues

Table 1: HERCL Commands

	Input and Output		**Stack Manipulation and Arithmetic**	
`i`	fetch INPUT to input buffer	`#`	PUSH new item to stack	$\ldots\ldots \mapsto \ldots\ldots x$
`s`	SCAN item from input buffer to stack	`!`	POP top item from stack	$\ldots\ldots x \mapsto \ldots\ldots$
`w`	WRITE from stack to output buffer	`c`	COPY top item on stack	$\ldots\ldots x \mapsto \ldots\ldots x, x$
`o`	flush OUTPUT buffer	`x`	SWAP top two items	$\ldots y, x \mapsto \ldots x, y$
	Registers and Memory	`y`	ROTATE top three items	$z, y, x \mapsto x, z, y$
`<`	GET value from register	`-`	NEGATE top item	$\ldots\ldots x \mapsto \ldots\ldots(-x)$
`>`	PUT value into register	`+`	ADD top two items	$\ldots y, x \mapsto \ldots(y+x)$
`^`	INCREMENT register	`*`	MULTIPLY top two items	$\ldots y, x \mapsto \ldots(y * x)$
`v`	DECREMENT register		**Mathematical Functions**	
`{`	LOAD from memory location	`r`	RECIPROCAL	$..x \to ..1/x$
`}`	STORE to memory location	`q`	SQUARE ROOT	$..x \to ..\sqrt{x}$
	Jump, Test, Branch and Logic	`e`	EXPONENTIAL	$..x \mapsto ..e^x$
`j`	JUMP to specified cell (subroutine)	`n`	(natural) LOGARITHM	$..x \mapsto ..\log_e(x)$
`|`	BAR line (RETURN on `.|` HALT on `8|`)	`a`	ARCSINE	$..x \mapsto ..\sin^{-1}(x)$
`=`	register is EQUAL to top of stack	`h`	TANH	$..x \mapsto ..\tanh(x)$
`g`	register GREATER than top of stack	`z`	ROUND to nearest integer	
`:`	if TRUE, branch FORWARD	`?`	push RANDOM value to stack	
`;`	if TRUE, branch BACK		**Double-Item Functions**	
`&`	logical AND	`%`	DIVIDE/MODULO	$..y, x \mapsto ..(y/x), (y \bmod x)$
`/`	logical OR	`t`	TRIG functions	$..\theta, r \mapsto ..r\sin\theta, r\cos\theta$
`~`	logical NOT	`p`	POLAR coords	$..y, x \mapsto ..\mathrm{atan2}(y,x), \sqrt{x^2+y^2}$

for a fixed number of rounds (typically, between 600 and 1000) with every round adding exactly one new image to the gallery.

Within each round, a new artist is evolved using hierarchical evolutionary re-combination, as illustrated in Figure 1 (see [19, 20] for further details). HERCL is a general-purpose evolutionary computation paradigm which has previously been applied to tasks such as classification [24], control [25], string processing [26] and line drawing [27]. It uses a simple imperative language with instructions for manipulating a stack, registers and memory, thus combining features from linear GP and stack-based GP. The full list of HERCL commands is given in Table 1. During evolution, HERCL code from previous or related tasks (in this case, the code that produced every previous image in the gallery) is kept in a library and made available as material for genetic recombination, so that the artist has the opportunity to explore variations on an earlier theme.

Note that the only role played by the human in this process is in searching for and selecting images of the desired subject, and scanning through the 600 to 1000 images produced, to select a subset of images that are considered to have the greatest visual appeal. Thus, the human acts purely as an initial patron and final curator of the artworks; the process of actually generating them is completely autonomous and untouched by human hands.

© Bernard Gagnon | (public domain) | © Ricardo Liberato | © Rich Niewiroski | © A.Savin

(a) Opera House (b) Eiffel Tower (c) Pyramids (d) Golden Gate (e) Saint Basil's

© Jérôme Blum | © Allard Schmidt | (public domain) | © DannyBoy7783 | © Hans Peter Schaefer

(g) Notre Dame (f) Machu Picchu (h) Taj Mahal (i) Angel Oak (j) Grand Canal

Fig. 2: Examples of real images for each of the 10 landmarks.

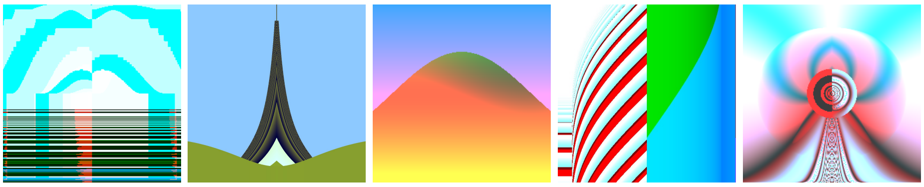

(a) Opera House (b) Eiffel Tower (c) Pyramids (d) Golden Gate (e) Saint Basil's

Fig. 3: Generated images from preliminary experiments using small critic network and no data augmentation.

4 Preliminary Experiments and Enhancements

For most of the experiments presented in [19], the real images covered the full range of CIFAR-10 categories [28], including planes, cars, birds, cats, deer, dogs, frogs, horses, ships and trucks. The generated images were intriguing, and exhibited a broad range of artistic styles; but, because the real images were so heterogenous, the generated images lacked a sense of coherence, and no image could really be attributed to, or said to "depict", any particular subject.

We therefore made a plan to extend that work with a new set of experiments in which photographs of a famous landmark would constitute the real images for each experimental run. We chose 10 landmarks, and in each case performed a Google image search and selected between 8 and 45 images which we felt would best serve as real images for that landmark (See Figure 2).

We first ran a set of five preliminary experiments, using the landmarks from Figure 2(a) to (e), and the same network structure used in [19], namely: a convolutional layer with 6 filters (5×5), a max pooling layer with stride 2, a second

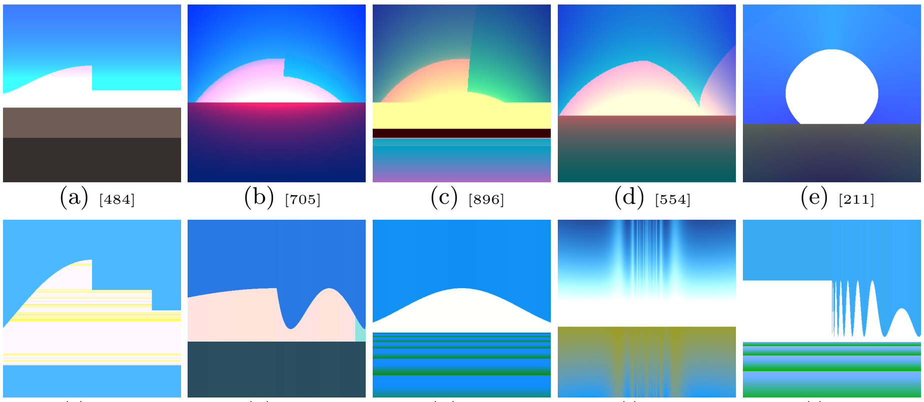

(a) [484] (b) [705] (c) [896] (d) [554] (e) [211]
(f) [445] (g) [444] (h) [108] (i) [348] (j) [124]

Fig. 4: Sydney Opera House

convolutional layer with 16 filters (5×5), another max pooling layer with stride 2, two fully connected layers of size 120 and 84, and two output units. Leaky ReLUs were used at all the hidden nodes, and softmax at the output. The real images were automatically converted to a resolution of 48×48, and all images generated by the artist were also rendered at a resolution of 48×48. The single image that we selected as most appealing from each of these five experiments are collected in Figure 3. Although these images do have some artistic merit, we felt that the system could perhaps be improved by increasing the number and size of the convolutional filters in the critic, and by employing data augmentation to compensate for the paucity of training images [29]. We also wanted to experiment with changing the resolution at which images were presented to the critic.

With this in mind, we ran two new sets of experiments, denoted Res48 and Res64. In both cases, the number of filters was increased in the first convolutional layer from 6 to 16, and in the second convolutional layer from 16 to 24. The size of the filters in the first layer was also increased from 5×5 to 7×7. For the Res48 experiments, images were generated at a resolution of 50×50 but cropped to 48×48 when fed to the critic. During training of the critic, data augmentation [29] was achieved by cropping both the real and generated images randomly in one of 9 different ways (stripping 0, 1 or 2 rows or columns from one end of the image and a complementary number from the other end). During evolution of the artist, exactly 1 row or column was stripped from all sides of the image and only the central 48×48 region was fed to the critic. For the Res64 experiments, images were generated at a resolution of 68×68 and cropped to 64×64 in one of 25 different ways for training of the critic, with exactly 2 rows and columns stripped from all sides of the image during evolution of the artist. The experiments were run until they had produced 1000 images for Res48, or 600 images for Res64 (in each case, about two weeks of computation on a single CPU).

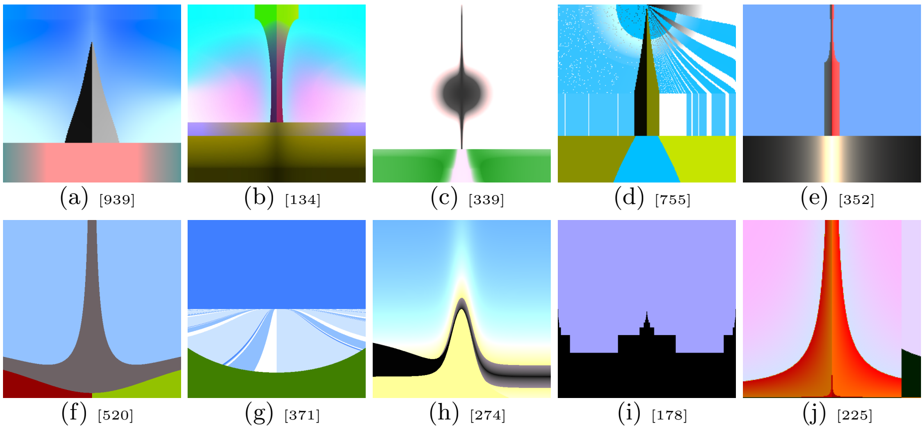

(a) [939] (b) [134] (c) [339] (d) [755] (e) [352]
(f) [520] (g) [371] (h) [274] (i) [178] (j) [225]

Fig. 5: Eiffel Tower

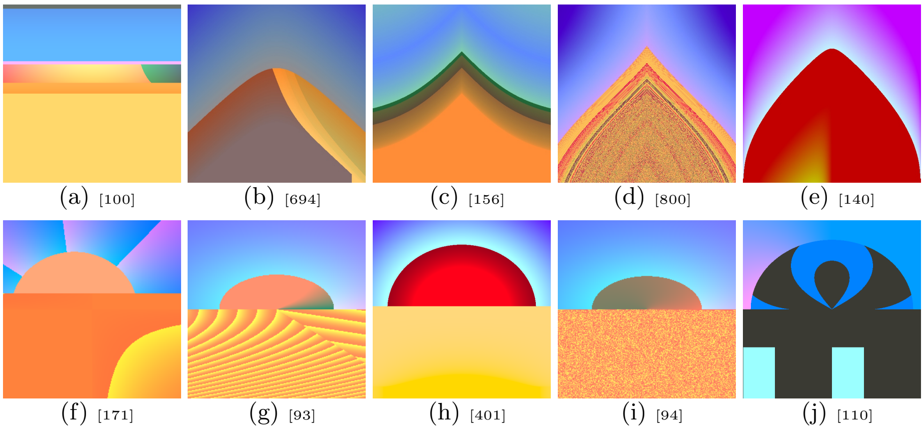

(a) [100] (b) [694] (c) [156] (d) [800] (e) [140]
(f) [171] (g) [93] (h) [401] (i) [94] (j) [110]

Fig. 6: Great Pyramid

5 Results, Artistic Styles and Analysis

Five selected images from each experiment are shown in Figures 4 to 13.[1] In each case, Subfigures (a) to (e) are from Res48 and (f) to (j) from Res64. The small numbers in square brackets indicate the round at which the image was generated.

5.1 Minimalism and Abstraction

Some of the generated images resemble minimalist art, where the subject is suggested with simple patters of colors and shapes. These include Figures 4(a, e, f, g), 5(a, f), 6(a, f), 7(a, b, f), 8(a, b), 9(f), 10(a, e, f, g, j), 11(a, g), 12(f, g) and 13(f).

Subtle yet important shading effects may be added to create an illusion of depth, as in Figures 5(j), 6(b, g, i), 7(c), 10(b, c, d, h), 12(i) and 13(b), or even to add a luminous quality, as in 4(b, c, d) and 6(h).

Convolutional neural networks with max-pooling layers and (Leaky) Rectified Linear Units generally make their classification based on an accumulation of visual features supporting the hypothesis of realness or fakeness. This semi-linear aggregation of evidence often enables the artist to fool the critic by providing some visual features in abundance while omitting others. This is evident in Figures 4(i), 5(g), 6(a) and 11(c) where the landscape is rendered but the building itself is missing; in 7(a, h, i) the vertical beams of the bridge are rendered but not its overall shape; Figure 11(e) includes the circular feature at the center of Notre Dame, and the general circular motif, but other features are absent. Figure 13(h) captures the color and texture (concentrated in the vertical direction) of buildings along the Grand Canal in Venice, but not their outline.

Components of the subject may be manipulated, such as in Figure 11(d) where the reflecting pool is visible but tilted at a weird angle, or 5(b) which gives the impression that a mischievous giant has uprooted the tower and stuck it back in the ground upside-down. Figures 8(a) and 9(a, j) have a slightly art deco look about them, whilst others such as 6(j), 8(j) and 13(a, e) seem to bear little resemblance to the underlying object, but instead make playful use of simple geometric forms in a limited range of colors, in the spirit of suprematism.

Figure 11(e) looks rather like a bird, totem pole or kite; yet, when we compare it to a real image of the Taj Mahal in 2(h), we see that the system has indeed extracted certain regularities of shape and color — but they are distinct from those that a human would normally focus on, thereby giving the work an abstract rather than figurative quality.

5.2 Colors and Shading

Art movements such as impressionism and fauvism rely in part on the propensity of the human visual system to respond to relative color rather than absolute color. This is often also the case for convolutional neural networks, because the early convolutional layers can learn to respond to differences in intensity among pixels in a local neighborhood, rather than absolute intensity. Our HERCL artist can exploit this feature of the LeNet critic by altering the colors of familiar objects. The Great Pyramid or the Eiffel Tower may become red; the Taj Mahal or Notre Dame may take on all the colors of the rainbow. In Figures 4(b, c, d), 5(a), 6(b, c, d, e, g, h, i) and 11(i) a halo effect is created to heighten the contrast (in one or more color channels) between the building and the sky. Some images such as 6(h), 8(b), 10(g), 11(j) and 13(j) are rendered in vibrant, perhaps even fauvist colors; others such as 7(c) and 10(h) are almost black-and-white in nature, but catch our attention through contrast and shading. Our HERCL artist seems to show a preference for vibrant reds and greens, which may be due in part to the encoding of the image in R,G,B space corresponding to the primary colors for mixing paints as well as the red, green and blue receptors of the human retina.

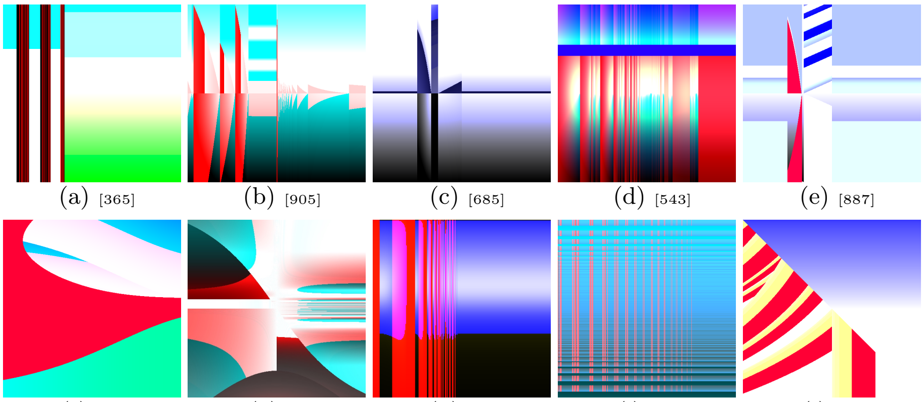

(a) [365] (b) [905] (c) [685] (d) [543] (e) [887]
(f) [476] (g) [80] (h) [501] (i) [573] (j) [221]

Fig. 7: Golden Gate Bridge

It would be interesting in future work to change the encoding to Y,U,V space and see whether a different palette of colors would then be preferred by the artist.

There are also cases where the artist has taken care to accurately reproduce the natural color. This typically occurs in situations where sections of the original photographs consist of a featureless expanse — such as the desert sands, or the sky over the Great Pyramid, Notre Dame or the Sydney Opera House.

Figure 8(b) provides an abstract rendition of Saint Basil's Cathedral in the form of colored ribbons. Comparing it to a real image in Figure 2(e), we see that it mimics both the shapes and colors of the original subject, but in new combinations. A central dome with a spire (depicted in red) appears to be flanked by two smaller, striped domes, with a couple of wispy white clouds, and black corners for the surrounding trees; blue, pink and yellow ribbons capture the colors of the blue dome, reddish towers and light colored base, while the green stripes reflect the green of the statue, the gables and the small tented roof.

5.3 Fractals and Texture

The selective pressure for images of low algorithmic complexity often entices the artist to employ self-similarity or fractals. Figures 8(c, d), 9(c, d, i), 11(d, j), 12(a, b, c, d, e) and 13(c, i) are recognizable as fractal art, and also exhibit some of the qualities of psychedelic art.

Figure 7(i) uses fractal geometry to tease the mind of the observer in a manner similar to the works of M.C. Escher. When we look at this image (best viewed in color) it is obvious that the red beams are reflected in the water; yet, we cannot quite discern where the sky meets the sea, or where the real beam ends and the reflection begins. Thus, the eye of the observer is led around in endless circles, searching for elusive boundaries that never quite resolve themselves.

Figures 4(h, j) and 10(d) use horizontal self-similarity to simulate topography or waves and ripples on the water, while vertical self-similarity is used in 4(i),

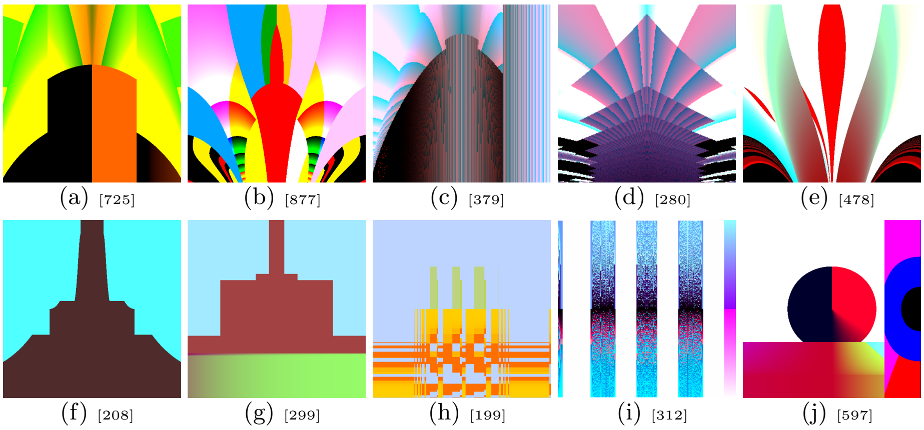

Fig. 8: Saint Basil's Cathedral (Sobor Vasiliya)

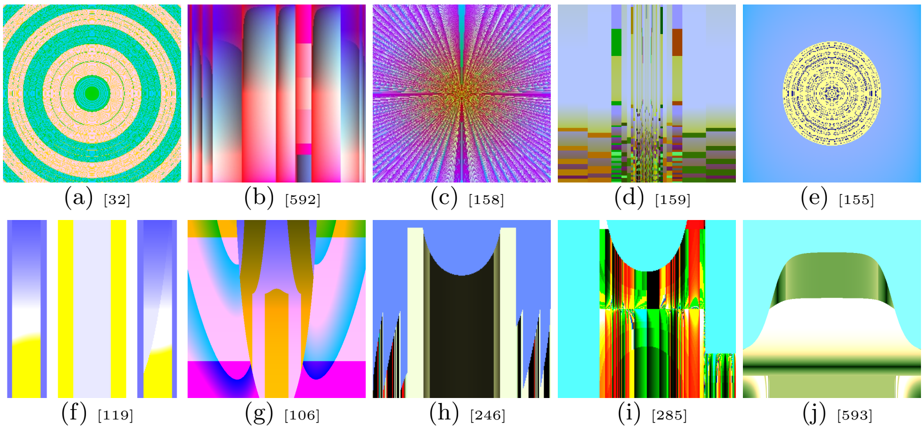

Fig. 9: Notre Dame de Paris

7(d, h, i) and 13(h) to generate beams or reeds growing thinner as they recede toward the horizon (or its vertical equivalent).

Fractals are used in Figures 8(c) and 8(d) to create the impression of a psychedelic tower. Circular self-similarity is used in Figures 9(a) and 9(e) to mimic the centerpiece of Notre Dame Cathedral, while a distinctive rectangular fractal structure is employed in 9(d) to simulate its overall shape. The fractal imagery in 9(i) gives the appearance of stained glass or glass art.

All the fractal images can be blown up to the size of wall paintings, with fine details emerging at the new scale (such as the delicate structure of the wings in Figure 12(e)). In parts of the images where the function defining the R,G,B values becomes extremely sensitive to changes in x and y coordinates, a seemingly random dot pattern can emerge, creating a form of pointillism. This phenomenon produces an exotic texture in Figure 8(i) and a "stardust" effect in

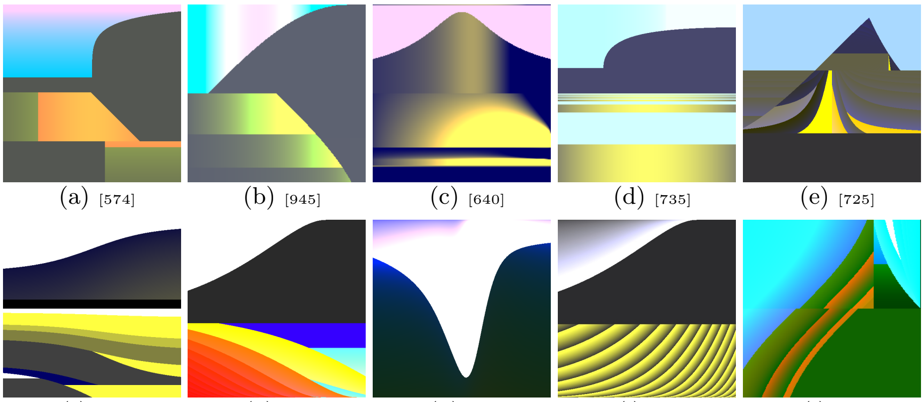

(a) [574] (b) [945] (c) [640] (d) [735] (e) [725]
(f) [536] (g) [295] (h) [408] (i) [351] (j) [459]

Fig. 10: Machu Picchu

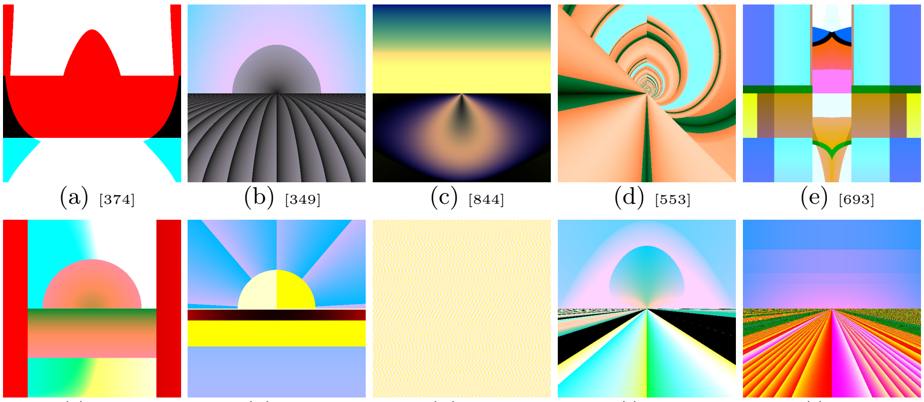

(a) [374] (b) [349] (c) [844] (d) [553] (e) [693]
(f) [157] (g) [475] (h) [239] (i) [198] (j) [350]

Fig. 11: Taj Mahal

5(d). In 6(d) a pattern that is regular near the outside of the pyramid appears to gradually break down to a random collection of dots as we move toward the center, thus achieving a kind of sandy texture. When Figures 6(d) and 6(i) are enlarged, we have the impression of looking out over a vast desert within which we can see individual grains of sand.

Figure 13(d) may at first appear to be a formless smattering of dots. However, with some imagination, we can think of it as conveying the general impression of a landscape reflected in the water, with a ridge along the waterline, some kind of indiscernible structures on the right, and open space to the left.

There is a tradition, dating back to Kazimir Malevich in 1915, of exhibiting an ostensibly blank canvas which contains no representational content in the usual sense, but instead invites the viewer to appreciate the texture of the paint or the detail of the brush strokes. Very occasionally, a similar phenomenon can be observed in our system; Figure 11(h) at first appears blank, but on closer inspection we can see that it contains fine details of texture. The LeNet critic, trained

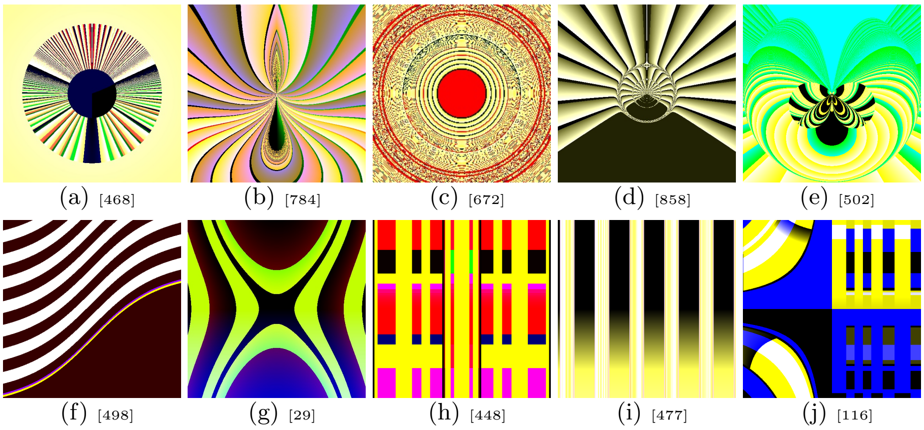

Fig. 12: Angel Oak Tree

to reject 238 previous images in the gallery, nevertheless assigns a probability of only 30% for this image to be synthetic.

5.4 Metaphor

There are several instances where images have been evolved to resemble one thing, but end up reminding us of something else. The Opera House becomes a rising moon in Figure 4(e); the Eiffel Tower looks like a cardboard cutout in 5(j); the Golden Gate Bridge can inspire a sailing rig in 7(e) or a collection of colored scarves hanging on a line in 7(j). Saint Basil's Cathedral prompts an image of a flower in 8(e); Notre Dame elicits an image resembling either a bed or a motor car in 9(j); the branches of an oak tree are morphed into a psychedelic insect in 12(e). Of course, our automated system has no knowledge of insects, ships, scarves, beds or flowers. It is possible that these metaphors emerge simply by chance. However, they may also reflect certain universal similarities that exist between various types of objects. If we suppose that Saint Basil's Cathedral was originally designed to resemble a flower, it may be that the wily artist has somehow brought to life the original conception behind the artifice.

One of the most intriguing metaphors appears in Figure 13(j) where a row of buildings along the Grand Canal in Venice have been rendered to look like a sailing ship. If we look closely at one of the original images in Figure 2(j) we can start to understand how the artist may have "mistaken" a triangular roof in the middle of the image for the prow of a ship, the buildings to its right for the hull, and the sloping roofs above for the sails. This example neatly illustrates the ability of the adversarial artist/critic system to surprise us by discovering connections and similarities that may otherwise have escaped our notice.

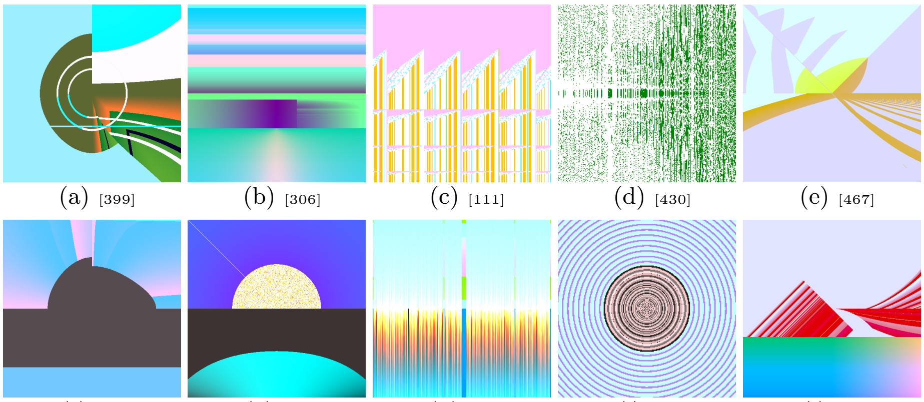

(a) [399] (b) [306] (c) [111] (d) [430] (e) [467]

(f) [458] (g) [572] (h) [198] (i) [195] (j) [214]

Fig. 13: Grand Canal in Venice

5.5 Repeated Substructures with Variations

In many cases we see substructures repeated within an image, but with minor variations. These include Figures 5(g), 7(j), 8(b, d, e), 9(b, d, g, i) and 12(b, h, j). As noted in [27, 30, 31] these imperfectly repeated substructures give the impression of having arisen organically through some natural process. In 13(j) the formless mass on the left appears to have been broken down and reconfigured into the sailing ship on the right; if we were to take the liberty of anthropomorphising the artist, we might be tempted to say that this work symbolises the use of raw materials to construct a fabricated object.

5.6 Genetype to Phenotype mapping

The HERCL code for the evolved images ranges in length from 10 to 223 characters, with a median of 73 characters (by convention, we treat the boundary between subroutines as a single character). The code length for the 100 selected images ranges from 36 to 162, with a median of 79 characters. As an example, the code (of length 73) for generating Figure 13(j) is shown in Table 2, along with equivalent pseudocode. Note that x runs from left (-1) to right $(+1)$, while y runs from top (-1) to bottom $(+1)$. Color intensities are output in the order blue, green, red; outputs less than 0 or greater than 1 are treated as 0 and 1, respectively. In order to avoid floating point exceptions, HERCL adheres to certain safety conventions, such as $\sin^{-1}(\alpha) = -\pi/2$ for $\alpha < -1$, which in this case ensures a uniform color for the sky.

6 Discussion

In general terms, the artistic quality across all landmarks (Figures 4 to 13) of the images from the Res48 experiments (Subfigures (a) to (e)) seems to be roughly

Table 2: HERCL code and Pseudocode for generating Figure 13(j)

<u>HERCL **code**</u>:

```
0[!qatcz]
1[capwwwo.]
2[%]
3[is.32#>sg:1j|c>xg:hp2j|+a{>cpa%.4338#p>g~<:0j|xww.88#wo]
```

<u>**Pseudocode**</u>:

scan (x, y) // $-1 \leq x \leq 1$, (upper) $-1 \leq y \leq 1$ (lower)
if $y \geq 0.32$ // water
  return $(\sqrt{y^2 + (\sin^{-1} y)^2},\ \mathrm{atan2}(y, \sin^{-1} y),\ x\,)$
else
  if $y > x$ // obstacle
    $u = \sin^{-1}(x + y)$
  else // ship
    $r = \sqrt{y^2 + \tanh(x)^2}$, $\theta = \mathrm{atan2}(y, \tanh(x))$
    $u = \sin^{-1}(\lfloor \theta / r \rfloor + (\theta \bmod r))$
  end
  $\phi = \frac{\pi}{4}(-1 + 2\,\mathrm{sgn}(u))$, $\rho = \sin^{-1}(\sqrt{2}u)$
  $z = \mathrm{atan2}((\phi \bmod \rho), 0.4338)$, $s = \sqrt{(\phi \bmod \rho)^2 + 0.4338^2}$
  if $s \leq z$ // sails
    return $(z, s, 0.88)$
  else // hull
    $v = \sin^{-1}(\sqrt{z}) \cos(\lfloor \phi / \rho \rfloor)$
    return $(v, \lfloor v \rceil, 0.88)$ // (blue, green, red)
  end
end

comparable with those from Res64 (Subfigures (f) to (j)). Subjectively, if we were asked to select the single most appealing image from each landmark, our choice would probably be: 4(c), 5(d), 6(d), 7(j), 8(b), 9(d), 10(e), 11(e), 12(e) and 13(j) — thus, 8 images from Res48 and only 2 from Res64. On the other hand, of these 8 images from Res48, 6 of them were generated later than the 600th round, which was the cutoff point for Res64. So, it would appear that Res48 is more likely to produce a high quality artwork within the computational budget of two weeks of CPU time; but, we cannot exclude the possibility that Res64 might produce superior images if it were run for longer than 600 rounds.

7 Conclusion

We have demonstrated the ability of an autonomous artist to create abstract art based on famous landmarks. The art is not constrained by human input, but instead emerges naturally from a tradeoff between the selective pressure for low algorithmic complexity and the imperative to fool the critic — thus giving the artist freedom to create images in its own intrinsic style. Nevertheless, we can recognize certain fundamental elements associated with influential art movements of the 19th and 20th Century including minimalism, impressionism, fauvism, pointillism and suprematism as well as psychedelic and fractal art.

The pressure for low algorithmic complexity often leads to self-similarity, fractals or repeated substructures which are pleasing to the eye. Our appreciation of these works suggests there may be some component of the human visual system which specifically responds to self-similarity. In future work, we hope to explore how such a mechanism might be modeled artificially, and integrated into both generative and discriminative convolutional neural networks.

Acknowledgment

Thanks to Jacob Soderlund, Darwin Vickers and Tanakrit Udomchoksakul for contributing code, and to Jeff Clune, Ken Stanley, Yoshua Bengio, Oliver Bown and Gary Greenfield for fruitful discussions. This research was undertaken with the support of Akin.com, as well as resources from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

References

1. Galanter, P.: Computational aesthetic evaluation: past and future, in: Computers and Creativity, Springer, 255-293 (2012)
2. McCormack, J., Bown, O., Dorin, A., McCabe, J., Monro, G. & Whitelaw, M.: Ten questions concerning generative computer art, Leonardo 47(2), 135-141 (2014)
3. Dawkins, R.: The Blind Watchmaker: why the evidence of evolution reveals a world without design, Norton, New York (1986)
4. Sims, K.: Artificial evolution for computer graphics, ACM Computer Graphics 25(4), 319-328 (1991)
5. Kowaliw, T., Dorin, A. & McCormack, J.: Promoting creative design in interactive evolutionary computation, IEEE Transactions on Evolutionary Computation 16(4), 523-536 (2012)
6. Secretan, J., Beato, N., D'Ambrosio, D.B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J.T. & Stanley, K.O.: Picbreeder: A case study in collaborative evolutionary exploration of design space, Evolutionary Computation 19(3), 373-403 (2011)
7. Machado, P., Correia, J., & Romero, J.: Expression-Based Evolution of Faces, EvoMusArt, Springer LNCS 7247, 188-199 (2012)
8. Nguyen, A., Yosinski, J. & Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 427-436 (2015)
9. Baluja, S., Pomerlau, D., & Todd, J.: Towards Automated Artificial Evolution for Computer-Generated Images, Connection Science 6(2), 325–354 (1994)
10. Ekárt, A., Sharma, D., & Chalakov, S.: Modelling Human Preference in Evolutionary Art, EvoMusArt, Springer LNCS 6625, 303-312 (2011)
11. Correia, J., Machado, P., Romero, J. & Carballal, A.: Feature Selection and Novelty in Computational Aesthetics, EvoMusArt, Springer LNCS 7834, 133-144 (2013)

12. Saunders, R. & Gero, J.S.: Artificial creativity: A synthetic approach to the study of creative behaviour, Computational and Cognitive Models of Creative Design V, Key Centre of Design Computing and Cognition, University of Sydney, 113-139 (2001)
13. Machado, P., Romero, J. & Manaris, B.: Experiments in computational aesthetics: An Iterative Approach to Stylistic Change in Evolutionary Art, The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music, Springer, 381-415 (2008)
14. Greenfield, G. & Machado, P.: Simulating Artist and Critic Dynamics — An Agent-based Application of an Evolutionary Art System, Proc. Int'l Joint Conf. on Computational Intelligence (IJCCI), Funchal, Madeira, Portugal, 190-197 (2009)
15. Li, Y. & Hu, C.: Aesthetic learning in an interactive evolutionary art system, European Conf. on the Applications of Evolutionary Computation, 301-310 (2010)
16. Correia, J., Machado, P., Romero, J., & Carballal, A.: Evolving Figurative Images Using Expression-Based Evolutionary Art, 4th International Conference on Computational Creativity (ICCC), 24–31 (2013)
17. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y.: Generative adversarial nets, Advances in Neural Information Processing Systems, 2672-2680 (2014)
18. Nguyen, A., Clune, J., Bengio, Y., Dosovitskiy, A. & Yosinski, J.: Plug & play generative networks: Conditional iterative generation of images in latent space, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 3510-3520 (2017)
19. Soderlund, J. & Blair, A.: Adversarial image generation using evolution and deep learning, IEEE Congress on Evolutionary Computation (2018)
20. Blair, A.: Learning the Caesar and Vigenere Cipher by hierarchical evolutionary re-combination, IEEE Congress on Evolutionary Computation, 605-612 (2013)
21. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P.: Gradient-based learning applied to document recognition, Proceedings of the IEEE 86(11), 2278-2323 (1998)
22. Rooke, S.: Eons of genetically evolved algorithmic images, in: Creative evolutionary systems, P.J. Bentley & D.W. Corne, Eds., Morgan Kauffmann, 339-365 (2002)
23. Kingma, D.P. & Ba, J.L.: Adam: a method for stochastic optimization, International Conference on Learning Representations (ICLR), 1-15 (2015)
24. Blair, A.: Transgenic Evolution for Classification Tasks with HERCL, Australasian Conf. on Artificial Life and Computational Intelligence, LNAI 8955, 185-195 (2015)
25. Blair, A.: Incremental Evolution of HERCL Programs for Robust Control, Genetic and Evolutionary Computation Conference (GECCO) Companion, 27-28 (2014)
26. Soderlund, J.,Vickers, D. & Blair, A.: Parallel Hierarchical Evolution of String Library Functions, Parallel Problem Solving from Nature, LNCS 9921, 281-291 (2016)
27. Vickers, D., Soderlund, J. & Blair, A.: Co-Evolving Line Drawings with Hierarchical Evolution, Australasian Conference on Artificial Life and Computational Intelligence (ACALCI), Springer LNAI 10142, 39-49 (2017)
28. Krizhevsky, A.: Learning multiple layers of features from tiny images, Master's Thesis, Computer Science, University of Toronto (2009)
29. Krizhevsky, A., Sutskever, I. & Hinton, G.: ImageNet Classification with Deep Convolutional Neural Networks, Advances in Neural Information Processing Systems, 1097-1105 (2012)
30. Barnsley, M.: Fractal Image Compression, AK Peters, Natick, MA (1993)
31. Schmidhuber, J.: Low-complexity Art, Leonardo 30(2), 97-103 (1997)