

Research Article

Image Evolution Using 2D Power Spectra

Michael Gircys  and Brian J. Ross 

Brock University, Department of Computer Science, 1812 Sir Isaac Brock Way, St. Catharines, ON, Canada L2S 3A1

Correspondence should be addressed to Brian J. Ross; bross@brocku.ca

Received 2 August 2018; Revised 23 November 2018; Accepted 6 December 2018; Published 2 January 2019

Academic Editor: Colin Johnson

Copyright © 2019 Michael Gircys and Brian J. Ross. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Procedurally generated images and textures have been widely explored in evolutionary art. One active research direction in the field is the discovery of suitable heuristics for measuring perceived characteristics of evolved images. This is important in order to help influence the nature of evolved images and thereby evolve more meaningful and pleasing art. In this regard, particular challenges exist for quantifying aspects of style and shape. In an attempt to bridge the divide between computer vision and cognitive perception, we propose the use of measures related to image spatial frequencies. Based on existing research that uses power spectral density of spatial frequencies as an effective metric for image classification and retrieval, we posit that Fourier decomposition can be effective for guiding image evolution. We refine fitness measures based on Fourier analysis and spatial frequency and apply them within a genetic programming environment for image synthesis. We implement fitness strategies using 2D Fourier power spectra and phase, with the goal of evolving images that share spectral properties of supplied target images. Adaptations and extensions of the fitness strategies are considered for their utility in art systems. Experiments were conducted using a variety of greyscale and colour target images, spatial fitness criteria, and procedural texture languages. Results were promising, in that some target images were trivially evolved, while others were more challenging to characterize. We also observed that some evolved images which we found discordant and “uncomfortable” show a previously identified spectral phenomenon. Future research should further investigate this result, as it could extend the use of 2D power spectra in fitness evaluations to promote new aesthetic properties.

1. Introduction

1.1. Overview of Problem. Digital art brings to mind many wide and varying concepts and examples, with many digitally produced, original pieces finding their own acclaim [1, 2]. It is trivial for software to precisely replicate a digital image. On the other hand, we find it difficult to autonomously produce new images which share similar visual characteristics with images provided. Forming correct abstractions between digital data and their visual interpretations is an ongoing challenge covering many fields of study [3–6].

We focus on procedural textures, which are images generated with mathematical formulae and/or algorithms [7]. The terms “images” and “textures” are used interchangeably. Texture synthesis shows its use in applications ranging from interactive art systems [8], adaptive image filters [9], camouflage generation [10], and game asset generation [11] amongst others.

The ability to form minor alterations in these procedures allows us to easily make changes in a structured manner. However, it may not always be clear *a priori* how these changes will come to manifest. By combining together parts between the better performing generated images, we may gradually refine them and allow them to exceed the quality of any single prior image. With this process of evolutionary refinement, we are able to explore many similar images which can feature novel and creative variation. A technique to capture and replicate spatial properties would be of great benefit for improving these existing systems or expanding to new applications.

Evolutionary algorithms (EA)—and notably genetic programming (GP)—are able to nonexhaustively explore the space of possible images with little explicit understanding of how to affect high-level image changes [12–15]. Perhaps the most critical component in all EAs is the fitness measure, defining the metaheuristic which guides the search to optimal

solutions. With image synthesis, a bridge is needed to cross the divide from computer vision, information theory, and computational intelligence attributes we can evaluate from our rendering, to the psychological and cognitive understandings of perception.

With evo-art, we are often attempting to recreate characteristics of a target image, and not to precisely duplicate it. The idea of evolving near-matches, or “variations on a theme”, has been a goal in many previous applications [16–18]. Using an evolutionary approach, exact matches are possible for simple images, but become rather difficult for more complex targets.

In investigating the existing measures that can be computed from a rendered image, measures related to power spectral density appear to be promising. Estimates of power spectral density are based on the discrete Fourier transform of a signal, a measure of power across each component frequency. For 2D applications, a radial average of the 2D DFT coefficients with common polar distance (same spatial frequency) can be obtained for a more robust, abstract measure. A number of papers on image analysis/retrieval [4, 5, 19, 20] have been found which use this to more effectively classify images based on computationally tricky but perceptively obvious attributes (*i.e.*, Eastern versus Western art; Portrait versus Sketch versus Landscape). Despite this, little can be found relating to the use of power spectra for evolutionary art.

Power spectral density also plays a key role in spatial frequency theory. The theory purports that a human or animal visual cortex operates through coded signals in relation to observed spatial frequencies (in contrast to edge and line detection which can be prominently found in wavelets) [21–25]. An interesting adaptation of this research enables the identification of uncomfortable images through contrast and frequency analysis [3]. Power spectra of an image’s luminance were investigated, and certain frequency octaves were found to provide higher ratings of perceptual discomfort. We find numerous motivations toward the exploration of power spectral density as an art fitness measure, and promise in modelling perceptual spatial characteristics.

1.2. Goals. With spatial frequency being one of the more human-intuitive measures for shape and composition, and with the amount of existing research linking the measure to human perception, this paper shows its potential as a tool for guiding evolutionary textures. Our goal is to explore the use of these measures in evolutionary texture synthesis and evaluate their utility in production of digital evolutionary art. We consider our models of shape from a target image for use as a guide when evolving new images. It is hoped that by capturing and reproducing key spatial attributes of the image, we can see novel images with similar properties emerge in a creative exploration.

Our research presents a pair of milestones. Using genetic programming, we produce grayscale textures and explore the ability of Fourier-based fitness measures to replicate spatial properties of target images. The focus on grayscale images simplifies the texture formulae evolved, and permits experiments to concentrate on shape information. We then

explore the use of these measures for colour image synthesis. Most evo-art systems use colour, and so it is important to examine the applicability of our Fourier analyses to the colour domain. Doing so helps establish the utility of Fourier shape analysis as a tool for serious applications in evolutionary art.

1.3. Organization of Paper. The paper is organized as follows. Section 2 reviews the Fourier transform and its application toward 2D images. Section 3 discusses some of the important research literature of relevance to this paper, with a focus on evolutionary textures, and application of power spectral density measures. We outline the details of our experimental system in Section 4, and summarize the key findings of our initial experiments in Section 5. Later work with adaptations toward evolutionary art is discussed in Section 6. Conclusions are given in Section 7.

The paper presumes familiarity with genetic programming [14]. Further details of this research are in [26].

2. Background

2.1. Fourier Transform. The following briefly outlines some of the main technical details of Fourier analyses. A complete introduction is beyond the scope of this paper. We refer the reader to detailed discussions in [27–29].

Fourier analysis is a well-known tool which sees substantial use in signal processing applications [28]. The Fourier transform converts a signal with samples based on amplitude at points in time, to a representation which shows the power and phase of the signal’s constituent frequencies. The Fourier transform translates a signal into a sum of sinusoids, where the frequency of each periodic term relates to a component frequency found in the signal. The result of such a decomposition is typically encoded as a complex number for each frequency (see (1) to (3)).

$$f(t) = \sum_{n=-\infty}^{\infty} C_n e^{in\omega t} \quad (1)$$

$$C_n = \frac{1}{T} \int_0^T f(t) e^{-in\omega t} dt \quad (2)$$

$$= \frac{1}{2} (a_n - ib_n) \quad (3)$$

The real part of the coefficient (a_n) scales each term and may maintain its definition as the amplitude of the particular frequency. The additional imaginary component of the coefficient (b_n) can be used in conjunction with the real component to recover the phase of the frequency, as declared through the complex phase angle.

Adapting the Fourier transform to a 2D image can be done by applying the discrete Fourier transform (DFT) on each index of the first dimension, and then again along each row of the results. This gives us the amplitude and phase of how each frequency contributes to the total 2D signal. In applications with images, we often see most of the high-energy coefficients appear around the central positions and main axes of the shifted DFT [29], as seen in Figure 1.

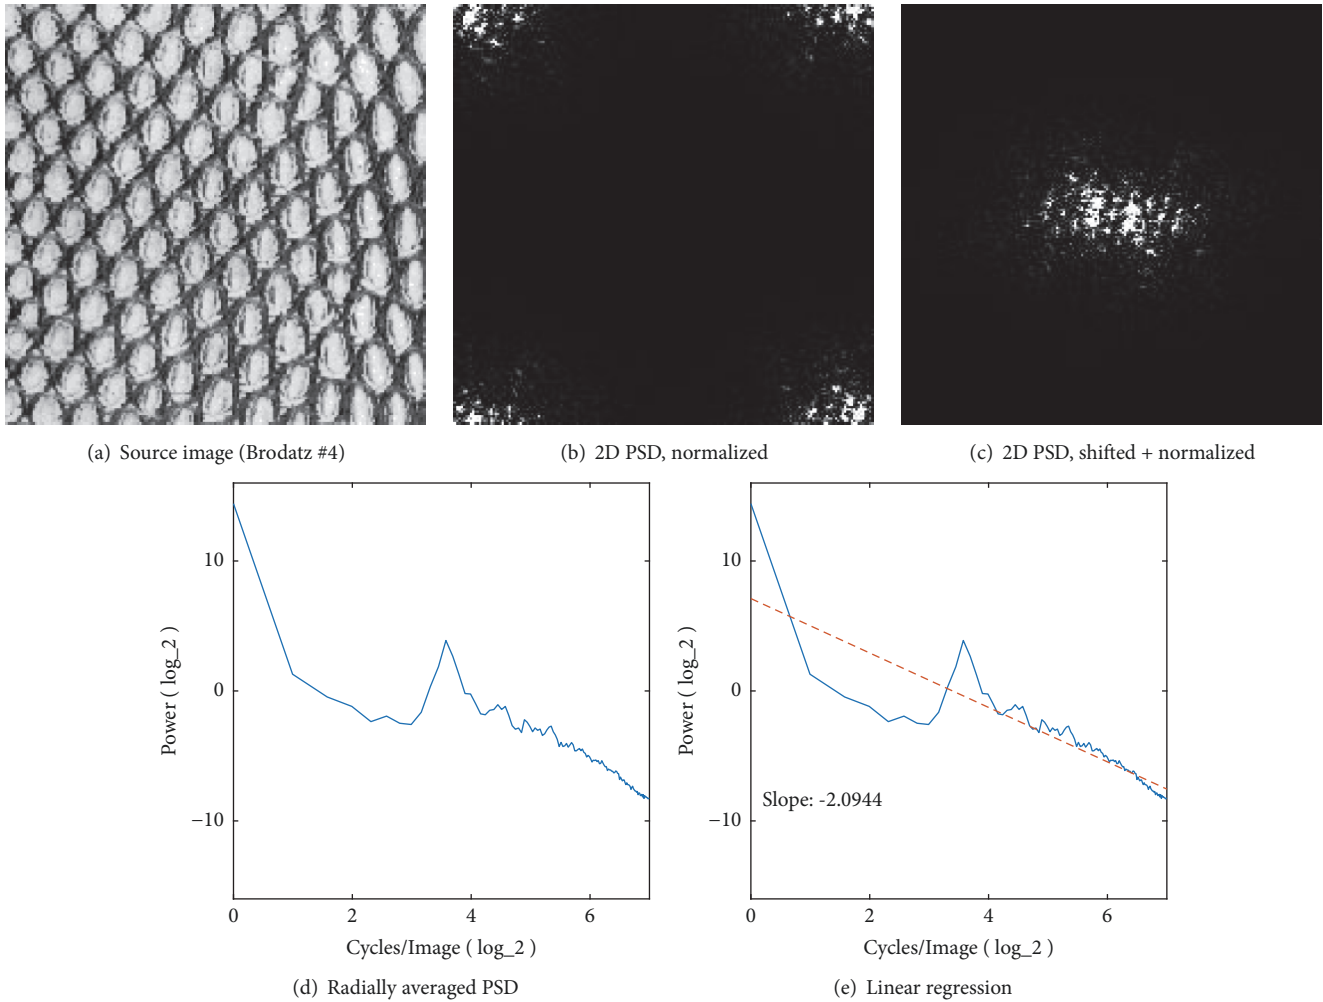


FIGURE 1: Power spectra pipeline. The source image in subfigure (a), undergoing 2D Fourier analysis, and its power spectral estimation shown in subfigure (b). As it provides for more interpretable charting and easier radial estimation, we “center” the coefficients by shifting them to diagonally opposite quadrants as seen in subfigure (c). We can then reduce dimensionality and produce useful aggregates by using radial averaging measures (d) and subsequent regressions (e).

Where the amplitude of an audio signal may have an intuitive correspondence with sound wave pressure, amplitudes for a 2D image will be measured in relation to their pixel intensity, or as is typically the case in colour images, the intensity across a particular colour channel.

While the Fourier transform can scale to higher dimensional signals, the use of DFT for colour textures is still potentially problematic [30]. In consideration of applying the DFT to colour channels in isolation, we should note that spatial properties are not necessarily clear from average intensity nor from inspection of individual colour channels. The related quaternion Fourier transform [31] might assist in this matter.

2.1.1. Power Spectral Density. The power spectral density (PSD), or power spectrum, is a measure of the power across the frequency domain of a signal. We can acquire an estimate of the PSD P_j at frequency j , by multiplying the Fourier

terms C_j by their complex conjugate $\overline{C_j}$ and scaling by the number of samples n to produce a periodogram [32]. Due to the simple, real-valued coefficients of our image signal, we can simplify this to normalizing and squaring the real part of the DFT, as in (4).

$$P_j = \left(\frac{C_j \overline{C_j}}{n^2} \right) \quad (4)$$

$$= \left(\frac{|C_j|}{n} \right)^2 \quad (5)$$

For a 2D signal, we will be interested in the radial average of this measure, requiring us to shift the quadrants of our estimate, and then interpreting the average in a polar coordinate system. An overview of the steps in our measurement pipeline is shown in Figure 1. Between the DFT and the radial averaging methods, the power spectral estimate measure has the benefit of being approximately equal across rotation, and

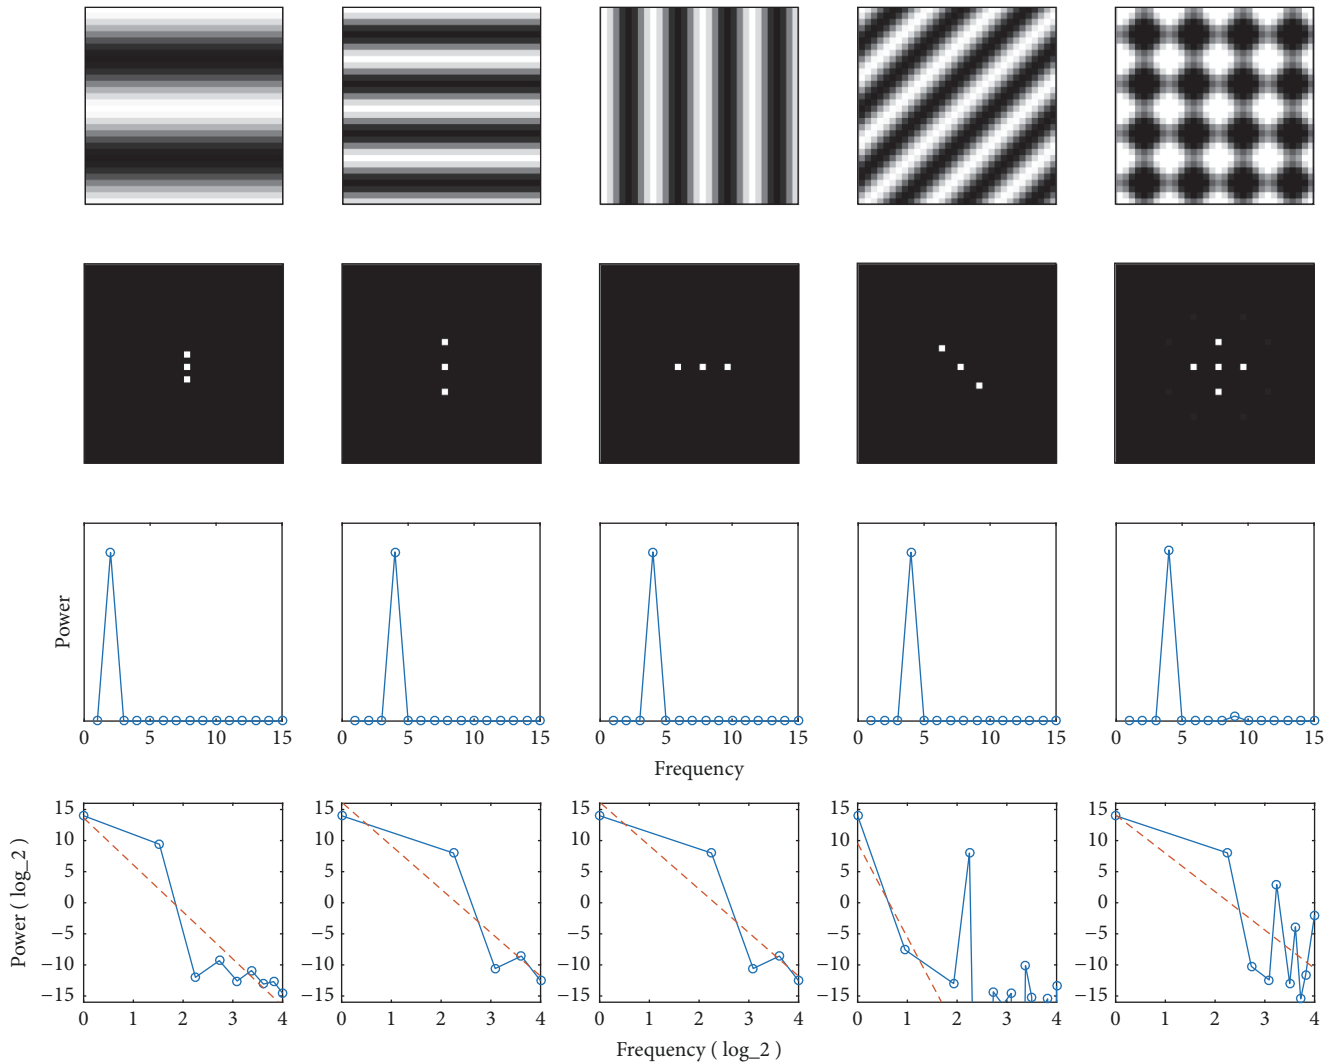


FIGURE 2: Power spectra interpretation and reconstruction. Rows from top to bottom: source image, shifted, and normalized constituent FFT power coefficients, radially averaged power spectra, and radially averaged power spectra plotted in a log-log scale (\log_2) with linear regression.

preserving shape across resolution. This measure relates to the contrast of luminance intensity, and we may also see a relation with image complexity. A further abstraction is to take a linear regression of the averaged power spectral density. While a display of the 2D power coefficients may more accurately represent the true power spectral density of a 2D signal, we find in some of the literature (*i.e.*, [4, 20]) that “power spectral density” and related terms often refer to the radial average or similar abstractions.

Figure 2 illustrates various representations of an image with a single component frequency. Shifting from the first to second column of the figure, we can see that lower frequency (those which have larger periods/cycles over greater areas of the image) is contained at the center of the shifted FFT power coefficient display. The first column shows a wave whose period is half of the canvas (input signal), and so the charted radially averaged power spectrum shows high power at a frequency of 2. As we move to the outer edge of the power coefficient display, we find the powers of increasing

frequency ranges being displayed. The fifth column faintly shows a suitable example of minor aliasing artefacts having both lower power and higher frequency as we move from the key frequencies toward the image edges. We can also observe that the orientation of the wave-like pattern in the top image corresponds to the angle (from center) of the coefficient responsible for the effect, while still maintaining a distance (from center) corresponding to the actual frequency. Observing the subsequently charted radially averaged power spectra plots, we can see that all have a high power at frequency 4. Finally, we can see the multiplicative combination of the two component frequencies in the last column, as a grid begins to form with both horizontal and vertical frequency, again reflected in the power coefficient display. The final row of the figure displays the radially averaged power spectrum in a log-log scale, to assist in showing the much larger 0^{th} coefficient, and the more subtle changes in the lower-powered high frequencies. However, in simple images, there may not always be power at every frequency. A problematic consequence of

this is that these frequencies cannot be charted in a log-log scale, and may affect the results of any regression, as is visible in the figure.

3. Literature Review

Although evolutionary algorithms have been applied to many forms of art over the years, we focus on literature involving the targeted evolution of procedural textures.

3.1. Spatial Measures. The need for measures permitting comparison of spatial properties tends to get resolved through one of two main concepts. Common approaches either extracted key features (and their positions) from a source or target image, or performed some type of frequency analysis. Many early attempts to capture spatial aspects for image database systems relied on basic algebraic and statistical measurements across intensity. The QBIC Project (which explored image querying through use of colour, texture, and shape measures) proposed spatial measures derived from capturing intensity areas, circularity, eccentricity, axis orientation, and algebraic central moment information [6].

A notable paper pertaining to image retrieval was published by Jacobs *et al.* [33], in which the proposed algorithm was capable of efficiently extracting the key coefficients from wavelet analysis. Extracted coefficients are limited to the K greatest absolute values, before being quantized and compared for mismatch. While the algorithm may have been intended for a retrieval system, the comparative abilities of the measure proved effective in guiding evolutionary systems. In [33], the set of coefficients were “truncated” by zeroing all but the top K greatest absolute value coefficients. Following this was a “quantization”, setting all nonzero components into their sign of $\{-1, +1\}$. The total error between images could then be found by summing of differences between each truncated, quantized coefficient position. This quantization scheme was found to be quite beneficial; despite the resulting loss of precision, as “the mere presence or absence of such features appears to have more discriminatory power for image querying than the features’ precise magnitudes” [33].

3.2. Evolutionary Textures. The use of evolutionary algorithms for texture synthesis was pioneered by Sims [13], and used interactive user guidance, which enabled a user to gradually manipulate sets of graphical shaders to produce images fitting a desired aesthetic.

An early attempt in the transition to unsupervised approaches came from Baluja *et al.* [15]. Simple topologies of artificial neural networks were used in an attempt to learn a user’s aesthetic preferences by training against user ratings and groups of raw pixel values. This approach saw some shortcomings, but highlighted the need for abstracted image measures to be used as guides. The idea of learning aesthetic preferences through neural networks has since been revisited with the inclusion of multiple abstracted image measures with some reported success [34].

A critical successor to Sims’ work was the Genshade system by Ibrahim [16]. Genshade introduced unsupervised,

automatic fitness evaluation of images as generated by evolved Renderman shaders. Various image analyses were compared between the evolved images and a provided target image. These measures were used in lieu of user input to guide the evolution of textures toward those showing similar visual characteristics of the targeted image.

The Gentropy system by Wiens and Ross [17] expanded upon the unsupervised approach of Genshade by providing additional image analysis measures, and use of a simple procedural texture language, in contrast to Genshade’s evolution of high-level Renderman shaders. A suite of image analyses were performed during fitness evaluation, which benefited with the use of island-model parallelism for maintaining diversity and accelerating the quality of evolved results. Gentropy was later enhanced in [35] by replacing island-model evolution with multiobjective evaluation, by treating the different image analysis tests as separate objectives for Pareto ranking.

Genshade [16] and Gentropy [17] employ the techniques from [33], where spatial features were compared via these extracted coefficients from wavelet measures. The technique appears to have been successfully adapted for use with texture synthesis. Results of wavelet analyses in both systems were positive, although a comprehensive investigation regarding the extent of their abilities was not undertaken.

More recently, there have been developments in using aesthetic modelling to guide image evolution [36–39]. Aesthetic modelling is a pioneering frontier for art and image analyses, and proposed models are not yet mature enough to be comprehensive theories of artistic beauty and aesthetics. Nevertheless, these efforts attempt to use higher-level image analyses as guides for evolution, which contrasts to the lower-level image processing used by systems like Genshade and Gentropy.

Recent work by Tanjil [40] uses ideas from deep learning to guide evolutionary image synthesis. A heuristic is proposed that enables activation nodes of a deep convolution neural network (trained for classification) to be identified for use by fitness evaluation. Using a set of images sharing desired visual features, the heuristic determines the activation nodes of the network most likely to be activated by the visual characteristics of interest. These nodes are then used as guides by fitness. A number of experiments showed that the genetic programming system was able to evolve images which shared desired properties of target images, such as shape and colour. Tanjil concludes that, as deep learning networks become better understood, they may be even more effectively exploited by evo-art systems.

While these and other systems attempt to capture spatial attributes, that was only a part of their purpose as a more generic art system. There was no extensive evaluation of their spatial guidance capabilities, and the use of Fourier analysis in texture synthesis or aesthetic modelling has been left largely unexplored.

Further examples and surveys of evolutionary art can be found in [1, 2], and contemporary research is published at the annual EvoMusArt conference (<http://www.evostar.org/>).

3.3. Limitations. Although the use of wavelet-based analysis showed effectiveness, alternative approaches are possible. One considered problem with a frequency analysis approach was in the inability to effectively handle images with multiple colour channels [30]. One potential solution to this was proposed through the use of the quaternion Fourier transform [31], which does not have a direct equivalence with wavelet analysis.

A criticism common to all types of frequency analysis remains in the fact that a perfect solution would exactly replicate the target image [41]. In evolutionary art, we never desire to make a reproduction of a given image. Rather, we only want to capture key characteristics of an image, and explore the landscape of possible solutions which are in some way similar. While fitness evaluations could be adjusted to prefer some amount of error, we found that there is often still sufficient challenge presented to our system outside of toy problems, permitting for novel solutions to emerge while we pursue higher numerical accuracy.

4. System Design

There are two key components which form the core of our experimental system. The first component is a library which could process an image to provide the power spectral density (PSD), regression, and other FFT related measures. The second, and largest, component is the evolutionary system which used genetic programming to evolve and synthesize procedural textures.

4.1. Power Spectral Density Measures. A number of PSD-related calculations were required for this research. For example, the 2D power coefficient matrices, the radially averaged power spectral density, and its linear regressions. MATLAB [42] (release 2016a) was used to assist with computation of power spectral density measures. MATLAB allowed us to generate native C code, which was integrated into the Java-based evolutionary system (Section 4.2) through use of the Java Native Interface (or commonly, JNI) framework.

For the experiments using regression measures of the radially averaged power spectral density, the regression was obtained first by converting the power measures to a log-log scaling, to better match the conventional practices seen in the literature. Charting of PSD throughout this paper uses \log_{10} scaling to remain consistent with other charted scales, though evaluations used for the various applicable experiments have used a \log_e scaling. For a linear regression, the slope measures should remain identical across log bases, though the offset will vary. Regressions were found by using MATLAB's `polyfit` function, which itself performs a least-squares error fit. While uncommon for natural images, some abstract images produced by our system were found to have no power at certain frequencies. To lessen the biased effects of these values from the regression, any infinite or invalid power measures were removed from the set of points considered during the regression.

We decided to forgo any image windowing functions prior to sending the image data through the DFT and PSD

measure pipelines. The use of a windowing function has been advised for nonregular signals, such as typical nonrepeating images, to reduce heavy artefacts in the decomposition. Specific window functions and parameters would be dependent on the expected signal. However, initial trials using windowing did not significantly impact our results, and so windowing was henceforth ignored.

In summary, tests found that our library produced results closely matching existing literature, and specifically those from Graham *et al.* [20].

4.2. Genetic Programming Engine. The evolutionary art system we used to generate textures is a custom extension of the ECJ system (version 23), a Java-based system for genetic programming and other techniques [43].

We used a genetic programming tree representation to evolve symbolic expressions for procedural textures. Much of our early experimentation focused on spatial attributes of an image. We found that grayscale textures were not only adequate, but were indeed preferable over the artistic colour texture renderings. To suitably represent this, GP individuals needed only a single tree to evaluate luminosity or intensity. Later experimentation expanded to colour textures, and we consequently expanded our individuals to hold 3 trees; one tree was used for each colour channel in the RGB colour-space.

The wall-clock run times for the system configured for basic grayscale textures were found to be approximately 45 minutes per run, when executed using a single thread of an AMD FX-8350 processor. In this configuration, multiple runs were evaluated concurrently. With the parallel nature of the system, we could see substantial reductions in single-run execution time if reconfigured to use multiple threads. The introduction of noise operators and RGB colour channels each increased runtime by factors of approximately 6 and 3, respectively. Coloured textures using noise language operators required an approximate average of 12 hours for completion of a run.

4.2.1. GP Parameters. Table 1 lists the GP parameters normally used in our experiments. Although most are standard in the literature [14], a few require explanation. Three variants of ephemeral random constants (ERCs) were included corresponding to orders of magnitude, and each of the ERC nodes are instantiated to random values within their respective ranges. The introduction of ephemeral value mutation allowed for the randomized constants to be slightly altered by 1%, which permitted for finer adjustments to the rendered image. The ERC mutation operator had been included at a probability of 10% and was responsible for a proportional decrease in likelihood to execute the crossover operator. So as to remove the possibility of losing the best found individual in a generation, we allowed elitism for the single best individual of a generation to be retained unaltered in the subsequent generation.

The termination criteria for a run were the completion of 100 generations. While “perfect” individuals had been produced for some simple compositional targets, this was

TABLE 1: Genetic programming engine parameters overview.

Parameter	Value
Runs	30
Generations	100
Population Size	1000
Elitism	1
Sum-of-Ranks Fitness	
Diversity Penalty, Initial	10
Diversity Penalty, Increment	10
Generation 0	
Builder	Ramped Half & Half (see [14])
New Node Depth	2...6
Grow Probability	50%
Reproductive Operators	
Crossover	70%
Mutation	20%
ERC Mutation	10%
Crossover Max Depth	17
Mutation Max Depth	17
Mutation New Node Depth	5
Selection Method	Tournament
Tournament Size	3

otherwise a difficult problem, where finding such a “perfect” solution was not typically expected.

4.2.2. Texture Languages. The GP language is in Table 2. Standard mathematical operators were used, as well as specialized texture generating primitives. Optimized Perlin and simplex noise generators have been borrowed from [44, 45] respectively. The fractalsum, turbulence, and marble noises have been based on the Perlin noise implementation as originally conceived. For these noise variants, coordinate scaling had been used to ensure noise is applied across the $[-1, 1]$ rendering window. Initial experiments in Section 5 excluded the spatial and noise operators.

4.3. Multiobjective Evaluation. Some problems permit us to evaluate solutions with a single measurement, for example, the overall error in a regression problem. However, there are problems where multiple criteria are necessary. These metrics can be independent, or can interact in complex, nonlinear ways. Reconciling such factors into a single metric score, for example, by a weighted sum, can be challenging to do effectively, and detrimental to search. The field of multiobjective optimization is concerned with problems such as these, in which multiple objectives are involved in defining the search criteria for a problem [46].

A popular scheme for scoring multiobjective problem spaces is Pareto ranking [47]. With Pareto, individuals are scored in relation to the others in the population. Unfortunately, Pareto ranking is not suitable for problems involving more than 3 objectives.

Our system uses the sum of ranks (or average rank) strategy, which was devised for multiobjective problems

involving a high number of objectives (termed “many-objective” problems) [48, 49]. Sum of ranks encourages solutions to perform well across all considered objectives. It is also effective for problems having a large number of objectives (unlike Pareto ranking). The sum of ranks approach has been found effective in evolutionary art applications [35, 39].

Table 3 illustrates the calculations for sum of ranks. After obtaining the raw measures (O_i) for each fitness objective, each measurement is separately ranked (R_i) relative to other individuals in the population. The rank scores are normalized (N_i) by dividing each R_i by the maximum rank value for that objective. The normalized ranks are summed for each individual, resulting in a fitness measure. The sum of ranks score denotes an individual’s relative performance of its objectives relative to the population at large. The final column Rank shows the relative fitness quality of each individual in the population. For example, individual #1 has the best score in each objective relative to the rest of the population, and thus has the best (lowest) sum of ranks. Individual #3 has an extremely poor score of 99 for objective 2. However, this raw score is converted to a rank of 5, and therefore does not unduly penalize the final ranking.

By using sum of ranks in our system, we are able to maintain a consistent diversity penalty scheme across all experiments. For individuals whose ranks in all objectives are identical, the second individual would have a penalty of 10 added to each of their ranks. Additional individuals found with the same scores as the first would incrementally receive an additional penalty of 10 rank points (the fourth common individual would receive a total of +30, and so on). These penalties are used to maintain genetic diversity in the population by penalizing identical results.

5. 2D Fourier Fitness Strategies

5.1. Simple Regression and Error. We first considered the error between FFT decomposition from evolved individuals and its target at a high level of abstraction. Beginning with the technique common in the literature (e.g., [4]), we considered a fitness scheme which measured the difference between slopes found through linear regression.

Measures of linearly regressed, radially averaged power spectra displayed some effectiveness previously with classification and retrieval. Consequently, evaluating fitness through this measure seemed like a promising start. Previous literature showed an improved ability to distinguish genre by incorporating this measure, and it was hoped that some spatial property capable of distinguishing these genres might emerge in our evolutionary synthesis.

In selection of a target set (Figure 3), we focused our efforts on aspects of spatial composition similarity. Though visually simple, the target images included basic compositions which might be used for evolutionary art.

Some concerns arose early into the process of constructing the linear regression module for our GP system. While much of the earlier explored work focused on evaluating natural images or complex art pieces, little investigation had been done into simple synthesized textures. In the process of

TABLE 2: Genetic programming engine base language overview.

Category	Arity	Display	Description	
Variables	0	X, Y	Texel rendering coordinates ($-1 \leq X, Y \leq 1$)	
		Rho	Polar coordinate; distance from $\{0, 0\}$	
		Phi	Polar coordinate; angle about $\{0, 0\}$ to X axis	
Ephemerals	0	E[1]	Ephemeral in range $[0, 1]$	
		E[10]	Ephemeral in range $[0, 10]$	
		E[100]	Ephemeral in range $[0, 100]$	
Math	1	-	Negation / sign change	
		abs	Absolute value / magnitude	
		floor	Floor; lesser or equal whole integer	
		ceil	Ceiling; greater or equal whole integer	
		sin, cos, tan	Periodic, trigonometric functions	
		sqrt	Square root	
		exp	e (Euler's number) raised to the operand	
		pow2, pow3	The operand raised to a fixed power of 2 or 3	
		log_E, log_10	Natural log, and log of base 10	
		2	+, -, *	Addition, subtraction, multiplication
			/	Safe division; a zero divisor returns zero
		3	max, min, avg	The greater, lesser, or mean of two operands
			pow	arg[0] raised to arg[1]
Conditionals	4	lerp	Linear interpolation between arg[0] and arg[1] based on normalized (clamped to $[0, 1]$) arg[2]	
		IfGT	If arg[0] > arg[1] then arg[2], else arg[3]	
Spatial	1	Circle	Gives 1.0 where $Rho \leq \arg[0]$, otherwise 0.0	
	3	Shift	arg[0] evaluated in rendering position shifted by arg[1] horizontally, and arg[2] vertically	
		Tile	arg[0] evaluated in rendering position scaled and offset for a $\arg[1] \times \arg[2]$ window tiling	
Noise	0	Simplex [†]	Simplex noise generator	
		Marble [†]	Marble noise (see [7])	
	1	FractalSum [†]	FractalSum/Smooth noise	
		Turbulence [†]	Turbulence noise	

[†]All noise functions include a variant symmetric about the X and Y axis. These variants would have a Sym prefix and function otherwise identical to the base function.

TABLE 3: Example of sum of ranks for a 3-objective problem. Lower objective scores and ranks are preferred. The maximum rank for each objective R_i used for normalization is in boldface.

#	Objectives			Rank			Normalized Rank			ΣN_i	Final Rank
	O_1	O_2	O_3	R_1	R_2	R_3	N_1	N_2	N_3		
1	1	1	3	1	1	1	0.25	0.2	0.33	0.78	1
2	2	2	4	2	2	2	0.5	0.4	0.67	1.57	2
3	2	99	3	2	5	1	0.5	1.0	0.33	1.83	3
4	4	4	4	3	3	2	0.75	0.6	0.67	2.02	4
5	6	7	5	4	4	3	1.0	0.8	1.0	2.8	5

charting the linearly averaged power spectra, and producing its regression, a transform into the log-log scale was required. Often, simple geometric images would result in frequencies with zero power. These anomalous frequencies needed to be removed, which could have an impact to the quality of regression.

Some example solutions for the slope results are shown in Figure 4. One positive aspect is that GP easily evolved images with a high degree of fitness to the targeted slopes.

The slope measure alone was insufficient in capturing any sufficient amount of spatial details. We found our GP system invariably converged to visually simple textures. The fitness criteria was too easily satisfied, and language biases were prevalent through our choice of simple mathematical operators. Unlike the use of regressed slope in image classification where it was applied to highly defined image sets (artwork, natural photographs, etc.), GP was capable of finding trivial solutions with the given slope criteria. Other experiments




FIGURE 3: Compositional target set.

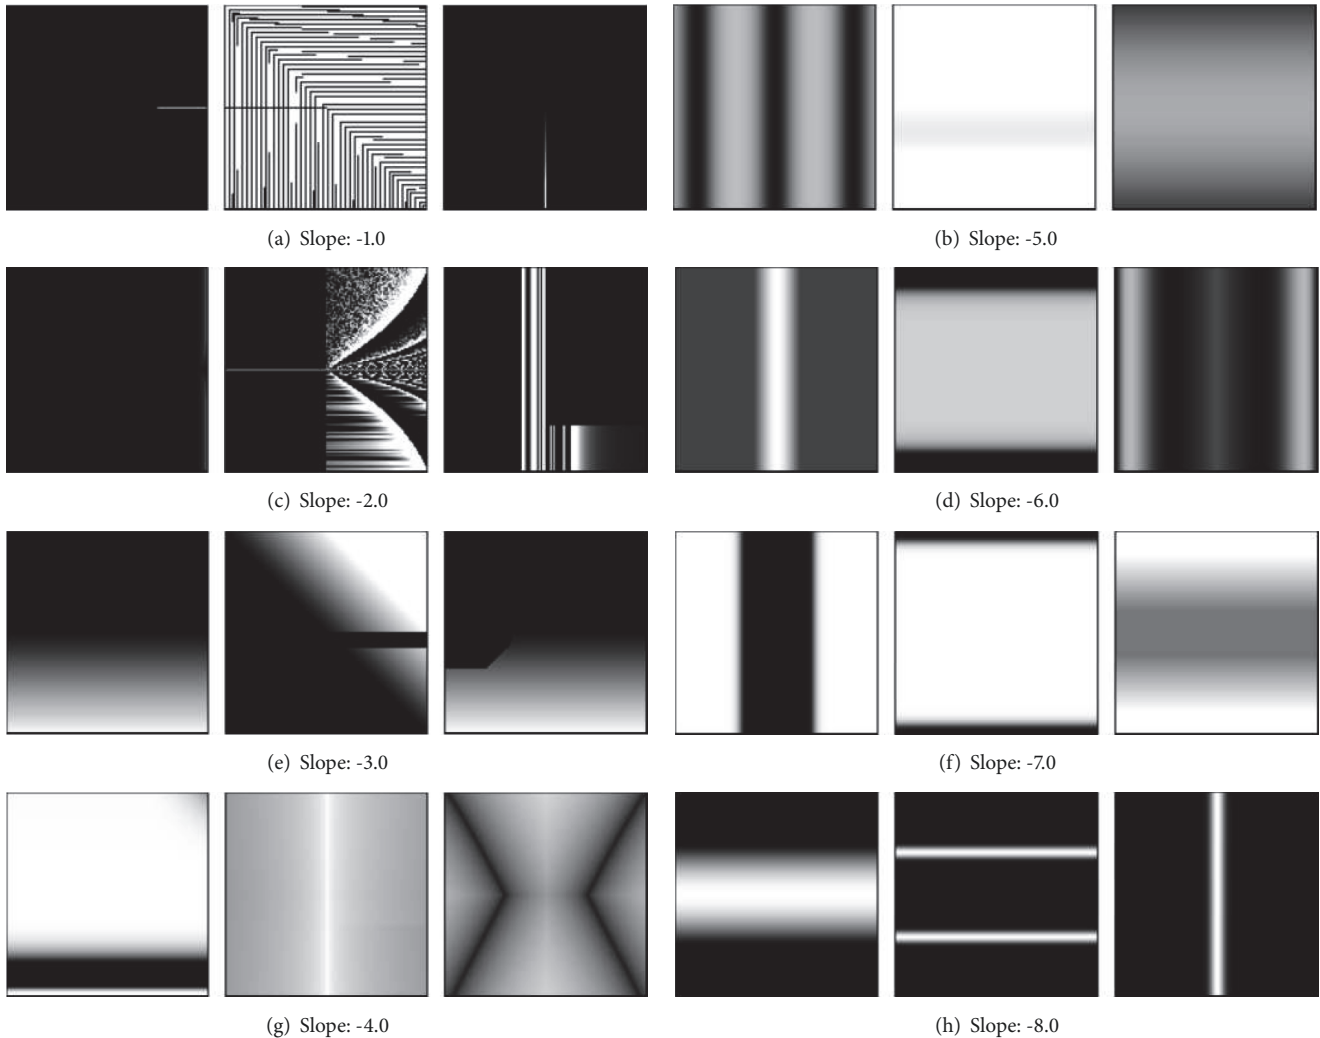


FIGURE 4: Regressed slope example evolved images. Target slopes were specified at regular integer intervals from -1.0 to -8.0 . Best candidates per run had mean $\epsilon < 1.0E^{-5}$, except for the last target, where it was found that errors greatly increased when target slope exceeded -7.0 . These initial runs relied on Cartesian coordinate variables (omitting the polar coordinate variables).

using power spectra regressions and similar basic measures were performed with only modest improvements to results (see [26]).

5.2. Filtering Relevant Coefficients. A promising strategy for coefficient isolation in frequency analysis was found by Jacobs *et al.* [33] using wavelets (See Section 3.1). There were a

few considerations to note before attempting similar schemes using Fourier transforms. A quantization to $\{-1, 0, +1\}$ was not as meaningful in the context of a Fourier transform, where power coefficients were strictly positive. Amplitude coefficients may have held negative values, but these could change sign when set with appropriate phase. We could truncate coefficients as per the paper, but the solution we

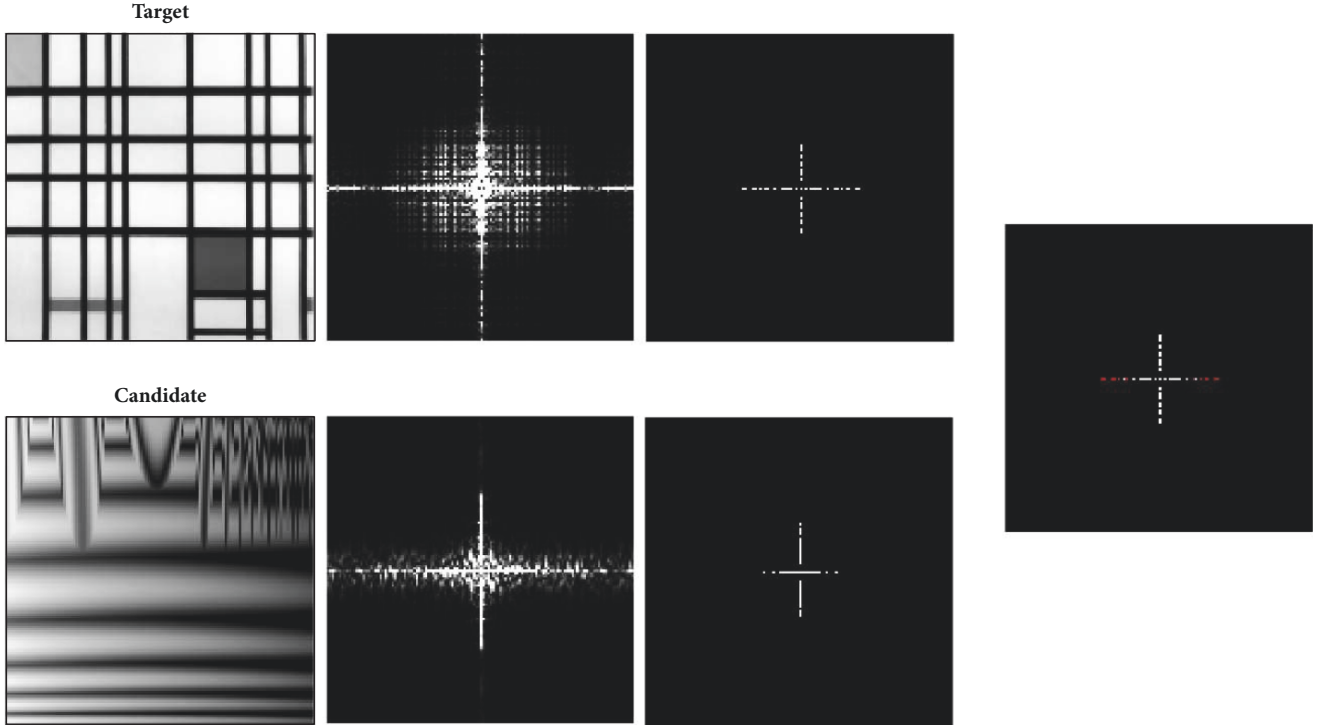


FIGURE 5: Truncation and quantization example. Target and candidate coefficient sets were reduced to the top $K = 50$ most powerful positions and then quantized to 0, 1, before checking for matches. The rightmost image shows the top K positions matching between target and candidate, with those absent from the target being displayed in red.

attempted instead quantized all remaining values in Boolean to 1. This effectively turned the score into a count of how many positions shared a top K coefficient between target and candidate. For a target and candidate of equal size, we ranked the target’s coefficient positions by their power, and truncated all but the top K . Each candidate could then undergo the same coefficient ranking process, and check its top K for a nonzero value in the corresponding location of the target’s truncated coefficients (see Figure 5).

While a wavelet decomposition would require further choices for wavelet type, decomposition type, and basis normalization schemes, Fourier compositions are constrained but simplified. A choice of K value was still required to determine the size of our coefficient truncation. Jacobs *et al.* found values of 40 to 60 performed well with their image retrieval data sets [33]. In our selection of a suitable K value, we considered possible reconstructions of the target images where power was removed from all but the top K positions. Prominent recreations began to form in the range of $K = [50, 150]$, where certain targets performed well with as low as $K = 10$.

5.3. Phase Refinement. A critical difference between the wavelet strategy of Jacobs *et al.* [33] using wavelets, and our adaptation with Fourier transforms, was the inherent removal of any spatial localization in our frequency analysis. When measuring coefficients, the index and position (the

radial angle of the coefficient from center) encouraged evolution of component frequencies with similar placement. However, this tended to overlook how these component frequencies should be offset and overlap. The other key aspect of a Fourier transform, the phase component, must therefore be considered. By reincorporating phase into our fitness scheme, we provided further constraints on the location of where the component frequencies crest. See Figure 6 for examples showing the effect of phase in Fourier reconstruction.

We adapted the Jacobs *et al.* approach—or, top K mismatch—and considered the difference of phase angle for those top K positions. Being mindful that phase error should wrap about 2π , the maximum difference in phase angle should be π . We normalized the phase error to $[0, 1]$ and squared it for each of the top positions. This error was then used to slightly penalize the top matching positions if they are out of phase.

We separated the phase error component to its own sum of ranks fitness objective, and applied a scaling factor on the phase to prioritize the more visually prominent (powerful) components. This is more formally defined in (6) and (7) and was also used for the next experiment.

$$Error_{power} = \sum_{i=1}^K \begin{cases} 0.0, & T_i \in C \\ 1.0, & T_i \notin C \end{cases} \quad (6)$$

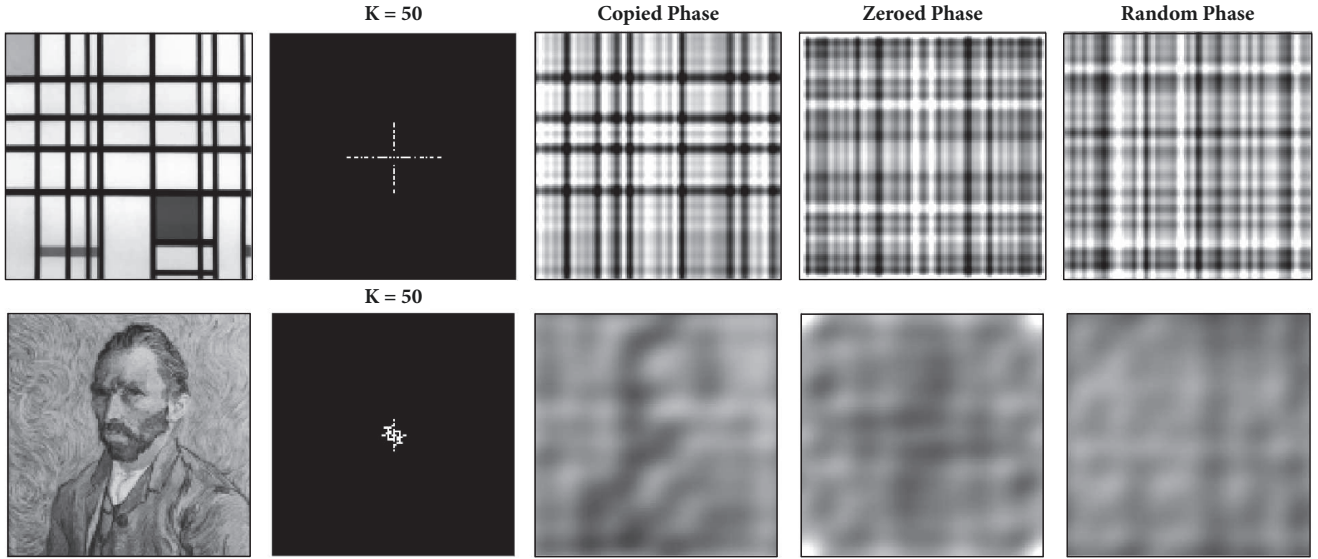


FIGURE 6: Reconstructing target images with varied phase. A pair of targets were chosen for reconstruction with a truncated set of Fourier amplitudes paired to different phase angle values. By zeroing power in all positions except the top $K = 50$ most powerful positions, we can see the most salient positions alongside the target in the second column. The third column is the inverted FFT reconstruction with power limited to these truncated positions. We also show reconstruction variants using the same truncated amplitude set, but with zeroed phase angles, or phase angles which have been produced randomly. This may adjust expectations for the types of images evolved when phase is not considered.

$$\begin{aligned}
 & \text{Error}_{\text{phase}} \\
 &= \sum_{i=1}^K \left\{ \begin{array}{ll} \left[\pi^{-1} \Delta(\theta(T, T_i), \theta(C, T_i)) \right]^2, & T_i \in C \\ 1.0, & T_i \notin C \end{array} \right. \quad (7) \\
 & \cdot \left(\frac{K-i+1}{K} \right)
 \end{aligned}$$

The equation assumes an $n \times n$ power coefficient set, where T and C are the truncated set of coefficient positions for the target and candidate as ordered by power. We have $\Phi(V, p)$ and $\theta(V, p)$ return the power and phase angle respectively of the coefficients (complex/vectors) in set V corresponding to coordinate p . With a slight abuse in notation, we denote the coordinates of the i^{th} ranked position (by power) of a coefficient set as S_i .

We show our key results in Figure 7. Using our measure, we were able to evolve images which show variations of their targets' key features. Similar regions of intensity can be seen for Composition_01, consistent horizontal stripes are produced for Composition_06, and vertical regions and gradients can be found in Composition_09 (some of which capture the finer details near its center). To have the regions of intensity seen in compositions 01 and 09 reproduced, proper capturing and recreation of phase information would be required. The low phase error seen for these targets (Table 4) is reflected in their visual similarity. The curves produced for the spiral target of Composition_10 are also quite interesting; the target was expected to be more difficult to satisfy, but we find variations of the key radial aspects are reliably recreated despite slightly elevated fitness error. Some notable examples produced have been highlighted in Figure 8.

Composition_06 (horizontal stripes) evolved candidates which scored well with our measure, and certainly captured the idea of horizontal stripes, but were not as uniform as seen elsewhere (see [26]). Despite closely matching the top $K = 10$ coefficients with its target, many evolved candidates also held large amounts of power in other coefficients. We found this was mitigated by adjusting K (at the cost of increasing outlier results), or trivialized by reducing the GP language. Particular difficulty was seen with Composition_07 (circle grid), but for different reasons. With this target, the produced solutions had high levels of error through our fitness measure. Our GP system allowed for the easy formation of unit circles and lines along the dimension axes, which makes for an underwhelming capture of the grid and circular aspects desired.

Extended runs terminating at 200 generations were attempted with little change to image quality. We can find further improvements on the targets with circular composition aspects by adjusting our GP language (Section 6.1.1). While certain targets may have performed better individually with various adjustments to the fitness measure (see [26]), the results from the above measure (shown in Figure 7) performed generally well across the majority of our target images.

6. More Advanced Artistic Explorations

Whereas Section 5 considered greyscale image synthesis, this section expands the scope of image evolution by considering more complex colour images. We first consider enhancements and extensions to our GP language which may better reflect some of the more full-featured languages used for

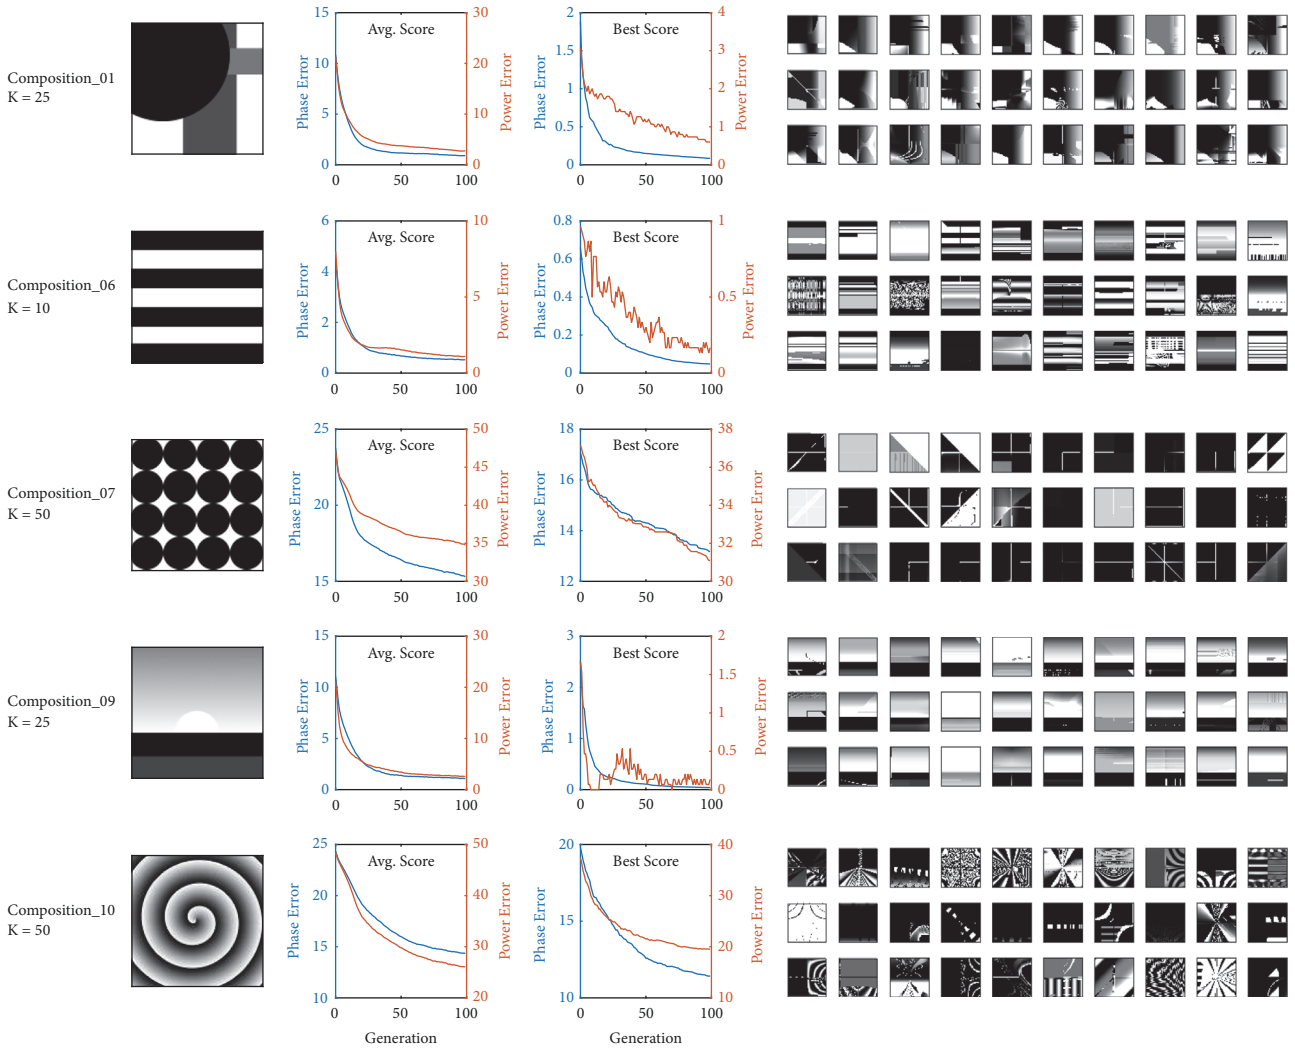


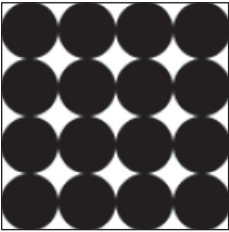




FIGURE 7: Compositional summary charts and examples. Each row of the figure captures the summary for a target over 30 runs. Leading with the target image, the next two columns show performance plots of the fitness measure (average over 30 runs). The leftmost plot displays the performance through the population average, where the rightmost plot shows the performance of the best individual of the generation. Aside the plots are the best candidate images produced at termination for each run.



FIGURE 8: Compositional experiment highlights. Images were produced using targets (from left to right): Composition_01, Composition_09, and Composition_10. Choice of K is outlined in Table 4. These examples show fair replication of compositional aspects, including placement of positions of intensity, contrasts and gradients, and shape characteristics.

TABLE 4: Compositional fitness summary table. Summaries for the remainder of the section were produced over 30 runs. For each target, a row is included for the mean and standard deviation for each fitness objective aggregate. Our experiment held two objectives, aiming to minimize error in power and phase coefficient matching. The row for “mean” shows the mean of terminal populations’ average fitness and the mean of terminal populations’ best found candidates across 30 runs. Maximum and expected error values are constrained by choice of K .

Target	K	Agg.	Power		Phase	
			Mean	Best	Mean	Best
	25	Mean	2.34	0.40	0.79	0.08
Composition_01		StdDev	0.68	0.81	0.18	0.05
	10	Mean	1.23	0.23	0.53	0.05
Composition_06		StdDev	0.38	0.43	0.17	0.05
	50	Mean	34.36	30.47	14.31	12.09
Composition_07		StdDev	3.26	4.44	1.16	1.46
	25	Mean	2.79	0.13	1.17	0.08
Composition_09		StdDev	0.78	0.51	0.36	0.13
	50	Mean	15.79	9.50	8.51	5.51
Composition_10		StdDev	4.82	4.66	2.89	3.04

evolutionary art applications. We then evaluate some possible multiobjective adaptations of our measures, and expand our capabilities from grayscale to coloured textures across multiple colour schemes. Finally, we present a brief discussion which corroborates a related measure in previously published research relating to computational aesthetics.

6.1. Language and Representation

6.1.1. Polar Coordinates, Geometric Operators. The first adjustment to our GP language was motivated by the poor performance observed when using targets which displayed strong radial attributes. We found that the inclusion of polar




FIGURE 9: Circle, grid, and offset language summary examples.

coordinate variables improved results for certain compositional targets (e.g., spirals) and some artistic genre targets. Some difficulty was still found with other targets using radial variations and repetitions. We therefore added a set of GP language operators well-suited for these target images.

With inspiration from the Gentropy system by Weins [50], we included the *circle* geometric operator (which returns 1.0 if the current texel is within the provided radius from the origin), along with the coordinate operators of *tile* and *shift*. The *circle* operator provided a simplified way for the candidate programs to show hard transitions about a radius, and the *tile* operator provided an easy way to create arbitrary $n \times m$ tilings.

Figure 9 shows a much-improved set of evolved candidate textures over our previous experiments. We see the error for these two targets decreases by $\sim 40\%$ in both objectives, and a 2-sample t-test provides at most $p < 0.0001$ across objectives and targets, suggesting fair statistical significance when considered with the reduced run count. The performance gains seen with these additional language operators is another promising sign for our fitness measure, and reinforces the importance for GP texture language adequacy.

6.1.2. Noise Generation. To help generate images having more visual complexity and interest, we included numerous noise generation operators (Section 4.2.2). With regard to error values, the introduction of the noise operators appears to be an improvement for most targets. We find minor but consistent reductions in both phase and power errors.

Figure 10 highlights some of the finer details in a pair of larger renderings using a target photograph of a flower.

6.1.3. Coordinate Variable Reduction. One final language experiment was performed by removing the X coordinate variable from the language set. It was expected that removing a fundamental coordinate variable would result in substantial difficulty for our system to produce results, and consequently, high error scores.

It is surprising to see that, despite the previous problems encountered while lacking the polar coordinate variables, there were few noted changes to performance. For the compositional target set, most targets performed only slightly better numerically with the inclusion of the X coordinate, and no statistical significance was found to favour either language set.

When we inspect the evolved textures a little more closely, there appears to be two main ways that our system and its textures adapted to the missing coordinate variable. Some candidates were able to glean sufficient positional information from the remaining coordinate variables: Y , ρ , and ϕ .

An alternative approach appears to largely forgo any direct positional information and instead builds upon layering multiple noise operators. We see this with the highlighted flower images in Figure 11, and a particularly interesting example of the Van Gogh target in Figure 12.

6.2. Colour. Here we considered the approach of evolving colour textures through separate evaluation of each colour channel, along with evaluation across average luminance. Further experimentation with HSL colour models, and other colour analyses can be found in [26].




FIGURE 10: Produced image highlights; noisy language; flower. The leftmost image shows the target, followed by a pair of notable evolved candidates.



FIGURE 11: Produced image highlights; noisy language, no X; flower. The leftmost image shows the target, followed by a pair of notable evolved candidates.

We maintained the selection of $K = 50$ as it produced suitable compositional results. To produce colour images, we evolved three GP trees per individual, corresponding to the RGB colour channels. With the increased tree count, and proportional increase in rendering complexity, we performed 9 runs per target. The system was then given 8 fitness objectives to optimize: the original grayscale power (Y) and phase (Y), colour power (R, G, B), colour phase (R, G, B).

6.2.1. $Y+RGB$ Colour Channels. As we found success with our existing measure on grayscale textures, we expanded upon this as a base. The placement and proportion of specific colours is guided using the same measurement technique across each individual RGB colour channel. Where a grayscale texture had two objectives (power and phase), our 4-channel ($Y+RGB$) colour image used $2 \times 4 = 8$ objectives. Each channel was evaluated similarly to a separate grayscale texture.

We maintained the use of a luminance channel evaluation as it was expected to further constrain the overall composition of the image. It was also hoped that the luminance

channel could capture some spatial information lost by assessing colour channels in isolation. We hypothesized that including this combination of luminance and colour channel objectives should reduce attempts to sacrifice any individual colour channel objective by incurring further penalties from mean luminance degradation. The NTSC (CCIR 601) method was used for conversion from colour (RGB) to grayscale:

$$Y = 0.299R + 0.587G + 0.114B \quad (8)$$

This provided a close approximation of colorimetric luminance from the nonlinear, gamma corrected RGB values.

The results in Figure 13 show that the control of colour through relative proportion and overlay of RGB channels, while basic and limited, is successful with certain targets. From the charting, we see similar sacrifices being made to the blue channel power error on target *Composition_15*. For *Composition_14-15*, the green channel, while still worse than when evolved in monochrome, sees some slight improvements. *Composition_13* sees an overall improvement to shape, where *Composition_01* remains consistent.




FIGURE 12: Produced image highlights; noisy language, no X; Van Gogh. On the left, we can see a snapshot of the candidate at every 20 generations. An evolutionary strategy has emerged which gradually applies layers and refines noise operators. The candidate is viewed atop the target image with partial transparency in the bottom left.

While we had hoped that the inclusion of a luminance channel would reduce the occurrence of sacrificing individual colour objectives, we occasionally see the opposite. There is now further pressure to sacrifice an objective if its channel is not contributing positively to the compositional shape as viewed through the lens of averaged luminance.

While overall colour distribution could be improved, we see increased performance when targets hold colour channels which can be replicated as grayscale targets individually. While considering the limitations, we are still able to replicate variations of shape and colour for a number of targets. Some highlights have been shown in Figures 14, 15, and 16.

6.3. Spatial Frequencies and Comfort. In the course of evolving the many candidate images with each target and

experiment set, we identified a number of evolved images which we found unpleasant or uncomfortable to view (see Figure 17). Previous research from Fernandez and Wilkins [3] found correlations between intensity level contrasts at certain spatial frequencies with increased levels of discomfort. We direct readers to their paper for an excellent example of the “uncomfortable property”.

The concept of spatial frequency denotes a cyclical nature across a measured space, such as the reoccurrence of Gabor and grating peaks along the width of an image. Our study is predicated over power coefficient positions directly relating to these spatial frequencies. While we found great utility in comparing spatial frequencies relative to image width, human perception requires consideration of an observers field of view. To better capture this, we can use calculations

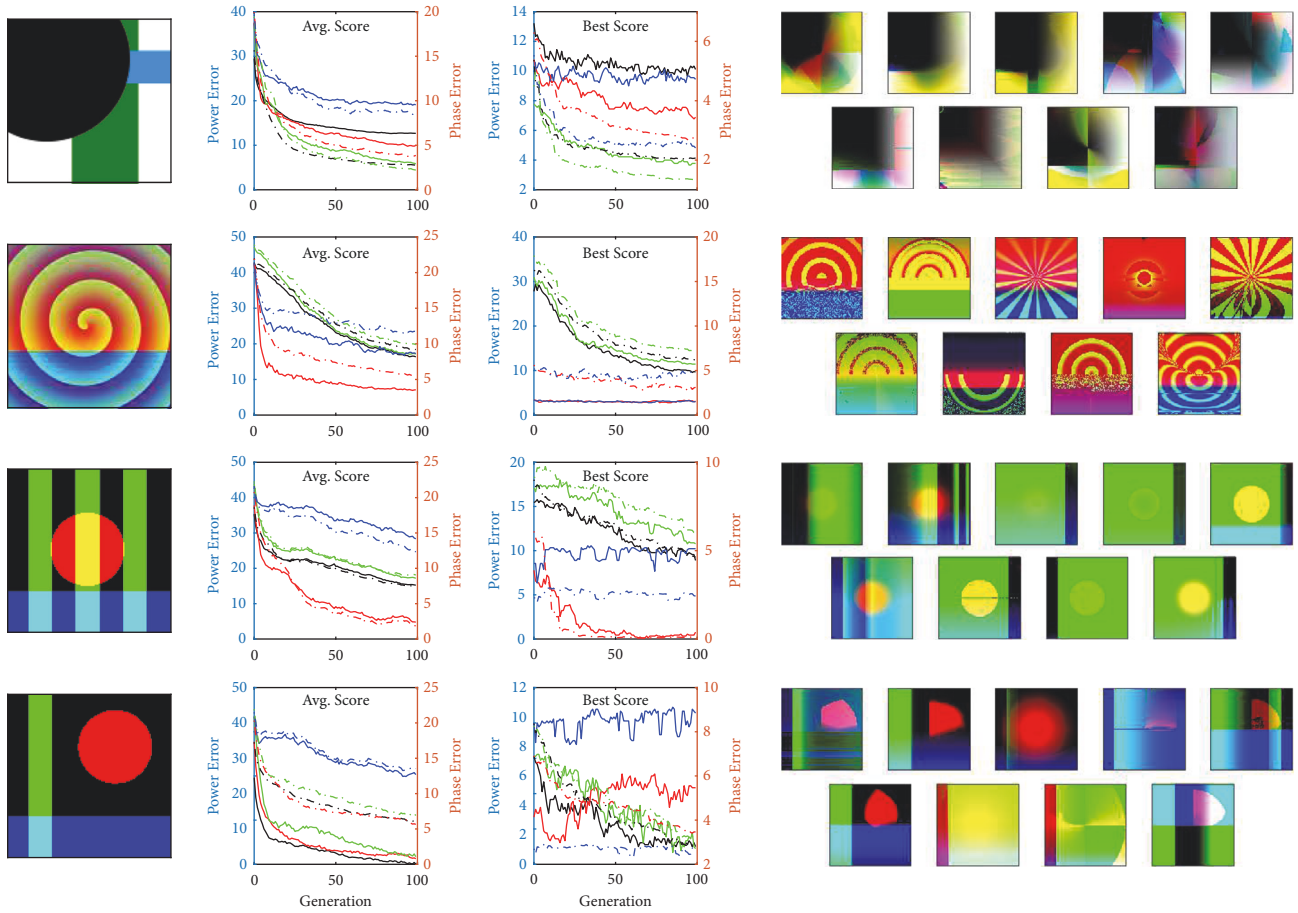


FIGURE 13: Colour experiment summary charts and examples. Power and phase errors for each of the red, green, and blue colour channels are plotted with their respective colour. Errors considered across the average luminance have been plotted in black. Power error is denoted with solid lines, where phase error uses a dashed line.




FIGURE 14: Colour experiment highlights; Composition_13; noise language. The leftmost image shows the target, followed by a pair of notable evolved candidates.

of visual angle – when paired with known viewing distance and image size – to compute a relative measure of angular spatial frequency. With spatial frequencies known in relation to image width, we can interpolate their corresponding visual angle when observed with known size and view distance.

Fernandez and Wilkins observed that images with increased amplitudes at a few octaves around 3 cycles per visual degree corresponded with higher reports of image discomfort. We explored numerous schemes in the previous sections to constrain and obtain specific spatial frequencies




FIGURE 15: Colour experiment highlights; flower; noise language. The leftmost image shows the target, followed by a notable evolved candidate.




FIGURE 16: Colour experiment highlights; flower; noise language, no *X*. The leftmost image shows the target, followed by a pair of notable evolved candidates.




FIGURE 17: Evolved images with uncomfortable spatial properties. Selection of images with uncomfortable aspects was performed with images sized to 12" side lengths at a viewing distance of 24" (identical angular spatial frequencies can be obtained when this page is viewed at a distance of 4.1", though we suspect that the eye strain induced from close proximity viewing will cause further undue discomfort).

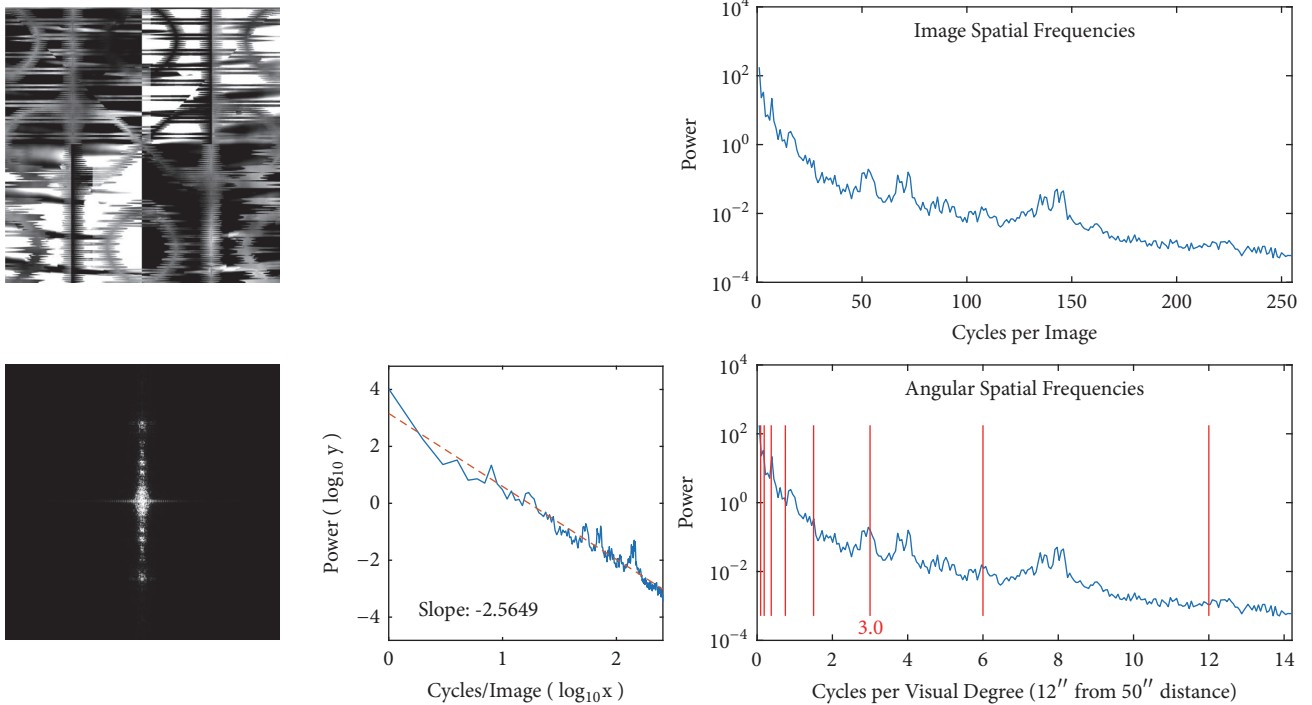


FIGURE 18: Angular spatial frequency analysis, distance variations. At the top left we have the analysed image, beneath which there are power spectra coefficients display and radially averaged power spectra. The top right graph plots power of the absolute spatial frequencies relative to image width, below which there are the spatial frequencies calculated relative to visual angle at a specific viewing distance. As recommended by Fernandez and Wilkins [3], octaves about 3 cycles/degree have been marked.

of a target image in our newly evolved candidates. With a direct relation between relative visual degree and absolute image spatial frequencies, we posit that these findings may be combined to the effect of a new aesthetic model.

Spatial frequency theory proposes that the human visual cortex operates through analysis of light receptor spatial frequencies [21, 22]. With supporting works finding sensitivity in animals to certain spatial frequency ranges [25, 51], it is not surprising to think that humans may also be more sensitive to contrasts at certain spatial frequencies. As seen in Figure 17, we can corroborate that intensity contrasts at certain visual frequencies are uncomfortable, discordant, and at times even painful. Figure 18 shows a frequency analysis of one of these evolved images. As can be seen, there is a peak in amplitude at the 3 cycles/degree frequency identified by Fernandez and Wilkins [3]. However, this measurement is dependent upon the viewing distance to the image, and this peak at 3 cycles/degree changes with different viewing distances.

With these findings, we identify a couple of limitations in using frequency for the analysis of uncomfortable images. The first, and least negotiable concern, holds that viewing size and distance must be considered before evolution. With interactive or hybrid fitness depending on user-evaluated thumbnails, large incongruities may appear between the rated thumbnails and full-size renderings.

There is another critical concern, though one we are now most capable of identifying and accommodating: naïve reduction to power within a range of frequencies can alter

an image to something unrecognisable. We have seen that core compositional information can be stored in 50 or so positions, as witnessed with our experiments in choice for truncation size, K . We can easily expect some of these critical frequencies to lay within the “ $3.0 \pm$ two octaves” range identified, and so a blanket frequency reduction should expect poor results with spatial similarity. If no other spatial attributes are sought in the evolved images, this penalty for power in the 3.0 angular spatial frequency range could provide a novel aesthetic measure for exploration. Some refinements will be needed otherwise. If provided target power spectra, we might propose an aesthetic objective which penalizes a surplus of power in these frequency ranges. From our observations above, we might also suggest a distribution of weights to provide harsher penalties when closer to the 3.0 cycles/degree mark.

Despite a number of concerns having been identified, our exploration with power spectra fitness measures has given us a tool to resolve some of them. We also suspect that beyond the correlation with discomfort and the given angular spatial frequency ranges, there may be a need to consider interactions with the phase of these frequencies and their harmonics. With further exploration in the future, novel aesthetic models can be developed from these findings.

7. Conclusion

2D power spectra can be an effective tool for guiding the evolutionary synthesis of images. By applying a 2D Fourier

analysis of a target image, key spatial characteristics can be extracted from it and used as a guide for the evolution of images that share these characteristics. Precise duplication of a target image is not desirable. Rather, by focussing on the major frequencies and their spatial orientations, the evolutionary art system is given enough freedom to “fill in the gaps” and generate interesting variations of images that have visual relationships to a target. Thus the approach acknowledges one of the strengths of evolutionary art, and evolution in general: the ability to generate creative and interesting solutions to problems.

Another unexpected result is the possible application of power spectra in identifying evolved images which have uncomfortable properties. A few example images show the spectral properties previously identified by Fernandez and Wilkins [3] in their study of uncomfortable art. Although more research on this topic is needed, there is the possibility of using such analyses within fitness strategies in order to avoid production of images with undesirable visual properties.

The success of the results shown in this paper depends upon two key factors. First, our coefficient reduction scheme proves effective in refining the search by simplifying the computational optimizations required in reproducing Fourier coefficients. Although further improvements and enhancements to this strategy are possible, our approach is generally effective for compositional targets and produced the results shown. Second, it is important that the procedural texture language used in the GP system has adequate power for producing images that conform to characteristics seen in the target image. The property of language adequacy and bias is well known in GP research. With our system, some target images are trivial to reproduce, where others are consistently difficult to handle with the basic procedural texture language. Improvements immediately arise when the language is supplemented with polar coordinates, noise generators, tiling operators, or other language features as needed by the target. On the other hand, some photographs we used as target images rarely yield successful outcomes, even with these additions. We hypothesize that our texture language remains incapable of easily generating images that match these targets. An enhanced texture language and coefficient reduction scheme may be warranted in these more challenging cases.

In summary, computer vision strategies such as spectral analysis continue to show wide success in applications involving image analysis, art classification, image retrieval, and other applications. These techniques should be given serious consideration in evolutionary art as well, in order to improve the quality and sophistication of machine-synthesized art.

Data Availability

The data used to support the findings of this study are available from the corresponding author upon request.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

This research was supported by NSERC Discovery Grant [RGPIN-2016-03653].

References

- [1] P. Bentley and D. W. Corne, *Creative Evolutionary Systems*, Morgan Kaufmann, 2002.
- [2] J. Romero and P. Machado, *The Art of Artificial Evolution*, Springer, Berlin, Heidelberg, 2008.
- [3] D. Fernandez and A. J. Wilkins, “Uncomfortable images in art and nature,” *Perception*, vol. 37, no. 7, pp. 1098–1113, 2008.
- [4] D. J. Graham, “Art statistics and visual processing: insights for picture coding,” in *Proceedings of the Picture Coding Symposium, PCS '09*, pp. 1–4, IEEE, May 2009.
- [5] D. Neumann and K. Gegenfurtner, “Image Retrieval and Perceptual Similarity,” *ACM Transactions on Applied Perception*, vol. 3, no. 1, pp. 31–47, 2006.
- [6] W. Niblack, R. Barber, W. Equitz et al., “Qbic project: querying images by content, using color, texture, and shape,” in *Proceedings of the IS and T/SPIE's Symposium on Electronic Imaging: Science and Technology*, pp. 173–187, International Society for Optics and Photonics, February 1993.
- [7] D. S. Ebert, *Texturing and Modeling: A Procedural Approach*, Morgan Kaufmann, 2003.
- [8] M. Hull and S. Colton, “Towards a general framework for program generation in creative domains,” in *Proceedings of the 4th International Joint Workshop on Computational Creativity, IJWCC '07*, pp. 137–144, June 2007.
- [9] M. Baniasadi and B. J. Ross, “Exploring non-photorealistic rendering with genetic programming,” *Genetic Programming and Evolvable Machines*, vol. 16, no. 2, pp. 211–239, 2015.
- [10] C. Reynolds, “Interactive evolution of camouflage,” *Artificial Life*, vol. 17, no. 2, pp. 123–136, 2011.
- [11] M. Hendrikx, S. Meijer, J. Van Der Velden, and A. Iosup, “Procedural content generation for games: a survey,” *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 9, no. 1, pp. 1–22, 2013.
- [12] J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, University of Michigan Press, 1975.
- [13] K. Sims, “Artificial evolution for computer graphics,” *ACM SIGGRAPH Computer Graphics*, vol. 25, no. 4, pp. 319–328, 1991.
- [14] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press, 1992.
- [15] S. Baluja, D. Pomerleau, and T. Jochem, “Towards automated artificial evolution for computer-generated images,” *Connection Science*, vol. 6, no. 2-3, pp. 325–354, 1994.
- [16] A. E. M. Ibrahim, *Genshade: an evolutionary approach to automatic and interactive procedural texture generation [Ph.D. thesis]*, Texas A and M University, 1998.
- [17] A. L. Wiens and B. J. Ross, “Gentropy: evolving 2D textures,” *Computers & Graphics*, vol. 26, no. 1, pp. 75–88, 2002.
- [18] P. Dahlstedt, “Turn-based evolution as a proposed implementation of artistic creative process,” in *Proceedings of the 2012 IEEE Congress on Evolutionary Computation, CEC '12*, pp. 1–7, IEEE, June 2012.
- [19] D. J. Graham and D. J. Field, “Variations in intensity statistics for representational and abstract art, and for art from the Eastern

- and Western hemispheres,” *Perception*, vol. 37, no. 9, pp. 1341–1352, 2008.
- [20] D. J. Graham and C. Redies, “Statistical regularities in art: Relations with visual coding and perception,” *Vision Research*, vol. 50, no. 16, pp. 1503–1509, 2010.
- [21] M. B. Sachs, J. Nachmias, and J. G. Robson, “Spatial-frequency channels in human vision,” *Journal of the Optical Society of America*, vol. 61, no. 9, pp. 1176–1186, 1971.
- [22] L. Maffei and A. Fiorentini, “The visual cortex as a spatial frequency analyser,” *Vision Research*, vol. 13, no. 7, pp. 1255–1267, 1973.
- [23] P. Vuilleumier, J. L. Armony, J. Driver, and R. J. Dolan, “Distinct spatial frequency sensitivities for processing faces and emotional expressions,” *Nature Neuroscience*, vol. 6, no. 6, pp. 624–631, 2003.
- [24] A. Fiorentini and N. Berardi, “Perceptual learning specific for orientation and spatial frequency,” *Nature*, vol. 287, no. 5777, pp. 43–44, 1980.
- [25] R. L. De Valois, D. G. Albrecht, and L. G. Thorell, “Spatial frequency selectivity of cells in macaque visual cortex,” *Vision Research*, vol. 22, no. 5, pp. 545–559, 1982.
- [26] M. Gircys, *Image evolution using 2d power spectra [M.S. thesis]*, Brock University, 2018.
- [27] J. M. Brayer, “Introduction to fourier transforms for image processing,” <https://www.cs.unm.edu/brayer/vision/fourier.html>.
- [28] P. Brémaud, *Mathematical Principles of Signal Processing: Fourier and Wavelet Analysis*, Springer Science and Business Media, 2013.
- [29] A. Rauh and G. R. Arce, “Sparse 2d fast fourier transform,” in *Proceedings of the 10th International Conference on Sampling Theory and Applications*, 2012.
- [30] S. J. Sangwine, “The problem of defining the fourier transform of a colour image,” in *Proceedings of the International Conference on Image Processing, ICIP ’98*, vol. 1, pp. 171–175, IEEE, October 1998.
- [31] V. R. Dubey, “Quaternion fourier transform for colour images,” *International Journal Computer Science and Information Technology*, vol. 5, no. 3, 2014.
- [32] National Semiconductor, “Power spectra estimation,” <http://www.dcs.warwick.ac.uk/feng/teaching/PowerSpectrum.pdf>.
- [33] C. E. Jacobs, A. Finkelstein, and D. H. Salesin, “Fast multiresolution image querying,” in *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, pp. 277–286, ACM, 1995.
- [34] Y. Li, C. Hu, L. L. Minku, and H. Zuo, “Learning aesthetic judgements in evolutionary art systems,” *Genetic Programming and Evolvable Machines*, vol. 14, no. 3, pp. 315–337, 2013.
- [35] B. J. Ross and H. Zhu, “Procedural texture evolution using multi-objective optimization,” *New Generation Computing*, vol. 22, no. 3, pp. 271–293, 2004.
- [36] E. Den Heijer and A. E. Eiben, “Investigating aesthetic measures for unsupervised evolutionary art,” *Swarm and Evolutionary Computation*, vol. 16, pp. 52–68, 2014.
- [37] P. Machado and A. Cardoso, “Computing aesthetics,” in *Proceedings of the XIVth Brazilian Symposium on AI*, pp. 239–249, Springer, Berlin, Germany, 1998.
- [38] P. Machado and A. Cardoso, “All the truth about NEvAr,” *Applied Intelligence*, vol. 16, no. 2, pp. 101–118, 2002.
- [39] B. J. Ross, W. Ralph, and H. Zong, “Evolutionary image synthesis using a model of aesthetics,” in *Proceedings of the 2006 IEEE Congress on Evolutionary Computation*, G. G. Yen, L. Wang, P. Bonissone, and S. M. Lucas, Eds., pp. 3832–3839, IEEE Press, Vancouver, Canada, July 2006.
- [40] F. R. Tanjil, *Deep learning concepts for evolutionary art [M.S. thesis]*, Brock University, 2018.
- [41] B. Julesz and T. Caelli, “On the limits of Fourier decompositions in visual texture perception,” *Perception*, vol. 8, no. 1, pp. 69–73, 1979.
- [42] The MathWorks, Inc., Mathworks: Matlab, <https://www.mathworks.com/>.
- [43] S. Luke, L. Panait, G. Balan et al., “Ecj: A java-based evolutionary computation research system,” 2006, <http://cs.gmu.edu/eclab/projects/ecj>.
- [44] Unported <http://creativecommons.org/licenses/by/3.0/>., “Riven, Perlinnoise: smooth/turbulent,” Creative Commons Attribution 3.0, <http://riven8192.blogspot.ca/2009/08/perlinnoise.html>.
- [45] S. Gustavson, “Simplex noise demystified,” <http://staffwww.itn.liu.se/stegu/simplexnoise/simplexnoise.pdf>.
- [46] C. A. C. Coello, G. B. Lamont, and D. A. Van Veldhuizen, *Evolutionary Algorithms for Solving Multi-Objective Problems*, Kluwer Academic, New York, NY, USA, 2nd edition, 2007.
- [47] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison Wesley, 1989.
- [48] P. J. Bentley and J. P. Wakefield, “Finding acceptable solutions in the pareto-optimal range using multiobjective genetic algorithms,” in *Soft Computing in Engineering Design and Manufacturing*, Springer, Berlin, Germany, 1997.
- [49] D. W. Corne and J. D. Knowles, “Techniques for highly multi-objective optimisation: Some nondominated points are better than others,” in *Proceedings of the 9th Annual Genetic and Evolutionary Computation Conference, GECCO ’07*, pp. 773–780, ACM Press, July 2007.
- [50] A. L. Wiens and B. J. Ross, “Gentropy: Evolutionary 2D texture generation,” in *Proceedings of the Late Breaking Papers at the 2000 Genetic and Evolutionary Computation Conference*, D. Whitley, Ed., pp. 418–424, July 2000.
- [51] N. P. Issa, C. Trepel, and M. P. Stryker, “Spatial frequency maps in cat visual cortex,” *The Journal of Neuroscience*, vol. 20, no. 22, pp. 8504–8514, 2000.

