

# Image Evolution Using 2D Power Spectra

*Michael Gircys*

Submitted in partial fulfilment  
of the requirements for the degree of

Master of Science

Department of Computer Science  
Brock University  
St. Catharines, Ontario

©*Michael Gircys*, 2018

# Abstract

Procedurally generated textures have seen use in many applications, are a high-interest topic when exploring evolutionary algorithms, and hold a central interest for digital art. However, there is an existing difficulty in finding suitable heuristics for measuring perceived qualities of an image. Particular difficulty remains for quantifying aspects of style and shape. In an attempt to bridge the divide between computer vision and cognitive perception, one set of proposed measures from previous studies relate to image spatial frequencies. Based on existing research which uses power spectral density of spatial frequencies as an effective metric for image classification and retrieval, we believe this measure and others based on Fourier decomposition may be effective for guiding evolutionary texture synthesis. We briefly compare some alternative means of using frequency analysis to guide evolution of shape and composition, and refine fitness measures based on Fourier analysis and spatial frequency. Our exploration has been conducted with the goals of improving intuition of these measures, evaluating the utility of these measures for image composition, and observing possible adaptations of their use in digital evolutionary art. Multiple evolutionary guidance schemes with consideration of the spatial frequencies' power spectra and phase have been evaluated across numerous targets with mixed results. We will display our exploration of power spectral density measures and their effectiveness as used for evolutionary algorithm fitness targets, particularly for basic compositional guidance in evolutionary art. We also observe and analyse a previously identified phenomenon of spatial properties which could lead to further consideration of visual comfort and aesthetics.

# Acknowledgements

To the following individuals, I would like to express particular appreciation for their continued support during my research:

- Brian Ross, for all your frequent insights and suggestions, and your endless patience and optimism. Each discussion spurred no shortage of great ideas to pursue.
- The members of my supervisory committee, Sheridan Houghten and Beatrice Ombuki-Berman, for your participation, advisement, and teachings.
- Cale Fairchild, for your patience and support while I sequestered your servers for a rendering farm.
- Fazle Tanjil, and Tyler Collins for their sharp second sets of eyes and assistance in ensuring program correctness.
- Samantha Laffin, for the insightful discussion and familiarization with spatial frequencies and visual angle.
- Family and friends, who kept me sane despite my long absences spent at a computer screen.
- Brock University, with note to the Department of Computer Science and the Department of Mathematics, for their generous computing resources, willing survey participants, funding, and facilities.

Without aid and encouragement from my family, friends, colleagues, and mentors, it is difficult to see this work being realized.

Thank you all for this opportunity!

# Contents

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>Introduction</b>                           | <b>1</b>  |
| 1.1      | Purpose . . . . .                             | 2         |
| 1.2      | Thesis Structure . . . . .                    | 3         |
| <b>2</b> | <b>Background</b>                             | <b>5</b>  |
| 2.1      | Procedural Textures . . . . .                 | 5         |
| 2.1.1    | Functional, Symbolic Expression . . . . .     | 6         |
| 2.1.2    | Noise Generation . . . . .                    | 7         |
| 2.1.3    | Colour Schemes . . . . .                      | 8         |
| 2.2      | Genetic Algorithms . . . . .                  | 9         |
| 2.2.1    | Representation . . . . .                      | 11        |
| 2.2.2    | Reproduction Operators . . . . .              | 13        |
| 2.2.3    | Evaluation . . . . .                          | 14        |
| 2.2.4    | Multi-Objective Schemes . . . . .             | 16        |
| 2.3      | Fourier Transform . . . . .                   | 17        |
| 2.3.1    | Power Spectral Density . . . . .              | 21        |
| 2.3.2    | Compare & Contrast: Wavelets . . . . .        | 24        |
| <b>3</b> | <b>Literature Review</b>                      | <b>26</b> |
| 3.1      | Evolutionary Textures . . . . .               | 26        |
| 3.2      | Power Spectral Density . . . . .              | 27        |
| 3.2.1    | Context: Vision . . . . .                     | 27        |
| 3.2.2    | Context: Classification & Retrieval . . . . . | 28        |
| 3.3      | Spatial Measures . . . . .                    | 28        |
| <b>4</b> | <b>System Design</b>                          | <b>30</b> |
| 4.1      | Power Spectral Density Measures . . . . .     | 30        |
| 4.2      | Genetic Programming Engine . . . . .          | 33        |

|          |  |            |
|----------|--|------------|
| 4.2.1    | Texture Languages . . . . .                                | 36         |
| <b>5</b> | <b>Initial Exploration</b>                                 | <b>39</b>  |
| 5.1      | Genre Images . . . . .                                     | 41         |
| 5.2      | Simple Regression & Error . . . . .                        | 43         |
| 5.2.1    | Linear Regression . . . . .                                | 43         |
| 5.2.2    | Radially Averaged Power Spectra . . . . .                  | 48         |
| 5.2.3    | Power Coefficients . . . . .                               | 52         |
| 5.3      | Filtering Relevant Coefficients . . . . .                  | 56         |
| 5.3.1    | Coefficient Extraction from Jacobs <i>et al.</i> . . . . . | 56         |
| 5.3.2    | The Importance of Phase . . . . .                          | 66         |
| 5.3.3    | Power Coefficient Error & Objective Separation . . . . .   | 73         |
| <b>6</b> | <b>Compositional Refinement</b>                            | <b>83</b>  |
| 6.1      | Compositional Targets . . . . .                            | 84         |
| 6.2      | Revisiting Previous Measures . . . . .                     | 84         |
| 6.3      | Adjusting Phase Error . . . . .                            | 92         |
| 6.4      | Survey Preparation . . . . .                               | 98         |
| <b>7</b> | <b>User Validation Survey</b>                              | <b>101</b> |
| 7.1      | Survey Design . . . . .                                    | 101        |
| 7.2      | Results & Analysis . . . . .                               | 103        |
| 7.3      | Conclusion . . . . .                                       | 104        |
| <b>8</b> | <b>Artistic Exploration</b>                                | <b>105</b> |
| 8.1      | Language & Representation . . . . .                        | 105        |
| 8.1.1    | Polar Coordinates . . . . .                                | 107        |
| 8.1.2    | Circle, Grid, & Offset . . . . .                           | 110        |
| 8.1.3    | Noise Generation . . . . .                                 | 112        |
| 8.1.4    | Coordinate Variable Reduction . . . . .                    | 116        |
| 8.2      | Colour . . . . .   | 121        |
| 8.2.1    | Colour Targets . . . . .                                   | 122        |
| 8.2.2    | Objectives Per Colour Channel . . . . .                    | 122        |
| 8.2.3    | Separating Shape & Colour . . . . .                        | 139        |
| 8.3      | Spatial Frequencies & Comfort . . . . .                    | 146        |

|                                 |            |
|---------------------------------|------------|
| <b>9 Conclusion</b>             | <b>154</b> |
| 9.1 Summary . . . . .           | 154        |
| 9.2 Future Work . . . . .       | 156        |
| <b>Bibliography</b>             | <b>166</b> |
| <b>Appendices</b>               | <b>167</b> |
| <b>A Target Analysis</b>        | <b>167</b> |
| <b>B Extended Results Data</b>  | <b>184</b> |
| <b>C User Survey &amp; Data</b> | <b>210</b> |

# List of Tables

|      |   |     |
|------|---|-----|
| 4.1  | MATLAB Code for 2D FFT Power and Phase . . . . .                | 31  |
| 4.2  | MATLAB Code for 2D Radial Average . . . . .                     | 31  |
| 4.3  | Genetic Programming Engine Parameters Overview . . . . .        | 35  |
| 4.4  | Genetic Programming Engine Base Language Overview . . . . .     | 37  |
| 4.5  | Genetic Programming Engine Extended Language Overview . . . . . | 38  |
| 5.1  | Overview of Experimental Variations . . . . .                   | 40  |
| 5.2  | Experiment E1 Summary Table (Slope Error) . . . . .             | 45  |
| 5.3  | Experiment E2 Summary Table (Slope + Offset) . . . . .          | 48  |
| 5.4  | Experiment E3 Summary Table (Curve Fitting) . . . . .           | 50  |
| 5.5  | Experiment E4 Summary Table (Coefficient Error) . . . . .       | 53  |
| 5.6  | Overview of Section 5.3 Experimental Variations . . . . .       | 56  |
| 5.7  | Experiment J1 Summary Table . . . . .                           | 59  |
| 5.8  | Experiment J2 Summary Table . . . . .                           | 63  |
| 5.9  | Experiment J3 Summary Table . . . . .                           | 70  |
| 5.10 | Experiment K1 Summary Table . . . . .                           | 74  |
| 5.11 | Experiment K2 Summary Table . . . . .                           | 78  |
| 5.12 | Experiment K3 Summary Table . . . . .                           | 79  |
| 6.1  | Overview of Section 6 Experimental Variations . . . . .         | 83  |
| 6.2  | Experiment P1 Summary Table; $K = 50$ . . . . .                 | 94  |
| 6.3  | Choice of $K$ for Compositional Targets . . . . .               | 99  |
| 8.1  | Overview of Section 8.1 Experimental Variations . . . . .       | 106 |
| 8.2  | Polar Language Summary Table . . . . .                          | 108 |
| 8.3  | Circle, Grid, Offset Language Summary Table . . . . .           | 111 |
| 8.4  | Noisy Language Summary Table . . . . .                          | 115 |
| 8.5  | Overview of Section 8.2 Experimental Variations . . . . .       | 121 |
| B.1  | Polar Language Summary Table; $K = 50$ . . . . .                | 206 |

|     |   |     |
|-----|---|-----|
| C.1 | User Survey Data & Partial Analysis . . . . . | 210 |
| C.2 | User Survey Raw Data By Question . . . . .    | 224 |

# List of Figures

|      |  |    |
|------|--|----|
| 2.1  | Rendering A Procedural Texture From A Function . . . . .                   | 6  |
| 2.2  | Rendering Window Bias Example . . . . .                                    | 7  |
| 2.3  | Fractal Noise Generation . . . . .   | 8  |
| 2.4  | Noise Generator Examples . . . . .   | 8  |
| 2.5  | Procedural Texture Separate Colour Channel Evaluation . . . . .            | 9  |
| 2.6  | HSL and RGB Colour Models Compared . . . . .                               | 10 |
| 2.7  | Genetic Algorithm and Genetic Programming Representation . . . . .         | 12 |
| 2.8  | Genetic Programming Reproduction Operators . . . . .                       | 13 |
| 2.9  | Fourier Synthesis in 1D . . . . .  | 18 |
| 2.10 | Power Spectra Pipeline . . . . .   | 21 |
| 2.11 | Power Spectra Interpretation & Reconstruction . . . . .                    | 23 |
|      |  |    |
| 5.1  | Initial Target Set . . . . .   | 42 |
| 5.2  | Regressed Slope Example Evolved Images . . . . .                           | 44 |
| 5.3  | Experiment E1 Summary Charts & Examples (Slope Error) . . . . .            | 46 |
| 5.4  | Experiment E2 Summary Charts & Examples (Slope + Offset) . . . . .         | 49 |
| 5.5  | Experiment E3 Summary Charts & Examples (Curve Fitting) . . . . .          | 51 |
| 5.6  | Experiment E4 Summary Charts & Examples (Coefficient Error) . . . . .      | 54 |
| 5.7  | Experiment E4 Target & Candidate Comparison . . . . .                      | 55 |
| 5.8  | Truncation and Quantization Example . . . . .                              | 58 |
| 5.9  | Experiment J1 Summary Charts & Examples . . . . .                          | 60 |
| 5.9  | Experiment J1 Summary Charts & Examples; $K = 150$ . . . . .               | 61 |
| 5.10 | Experiment J2 Summary Charts & Examples . . . . .                          | 64 |
| 5.10 | Experiment J2 Summary Charts & Examples; All Coefficients . . . . .        | 65 |
| 5.11 | Reconstructing Target Images With Varied Phase . . . . .                   | 67 |
| 5.12 | Reconstructing J1 Results With Target Phase . . . . .                      | 68 |
| 5.13 | Experiment J3 Summary Charts & Examples . . . . .                          | 71 |
| 5.13 | Experiment J3 Summary Charts & Examples; $K = 50$ ; Square Error . . . . . | 72 |

|      |  |     |
|------|--|-----|
| 5.14 | Experiment K1 Summary Charts & Examples . . . . .                      | 75  |
| 5.15 | Experiment K2 Summary Charts & Examples . . . . .                      | 77  |
| 5.16 | Experiment K3 Summary Charts & Examples . . . . .                      | 80  |
| 5.16 | Experiment K3 Summary Charts & Examples; $K = 150$ . . . . .           | 81  |
| 6.1  | Compositional Target Set . . . . .                                     | 85  |
| 6.2  | Experiment J3 Compositional Summary Charts & Examples . . . . .        | 86  |
| 6.2  | Experiment J3 Compositional Summary Charts & Examples; $K = 50$        | 87  |
| 6.2  | Experiment J3 Compositional Summary Charts & Examples; $K = 50$        | 88  |
| 6.3  | Experiment K3 Compositional Summary Charts & Examples . . . . .        | 89  |
| 6.3  | Experiment K3 Compositional Summary Charts & Examples; $K = 50$        | 90  |
| 6.3  | Experiment K3 Compositional Summary Charts & Examples; $K = 50$        | 91  |
| 6.4  | Experiment J3 Compositional Examples; $K = 10$ , Composition_04-05     | 93  |
| 6.5  | Experiment P1 Compositional Summary Charts & Examples . . . . .        | 95  |
| 6.5  | Experiment P1 Compositional Summary Charts & Examples; $K = 50$        | 96  |
| 6.5  | Experiment P1 Compositional Summary Charts & Examples; $K = 50$        | 97  |
| 6.6  | Polar Coordinate Language Summary Examples; $K$ Per Target . . . . .   | 100 |
| 7.1  | Example Survey Question . . . . .                                      | 102 |
| 7.2  | User Survey Compositional Targets . . . . .                            | 103 |
| 8.1  | Polar Coordinate System . . . . .                                      | 107 |
| 8.2  | Polar Coordinate Language Genre Targets Summary Charts & Examples      | 109 |
| 8.3  | Circle, Grid, Offset Language Summary Examples . . . . .               | 111 |
| 8.4  | Noisy Language Summary Examples; Compositional Targets . . . . .       | 113 |
| 8.5  | Noisy Language Summary Charts & Examples; Genre Targets . . . . .      | 114 |
| 8.6  | Produced Image Highlights; Noisy Language; Flower . . . . .            | 116 |
| 8.7  | Produced Image Highlights; Noisy Language, No $X$ ; Flower . . . . .   | 117 |
| 8.8  | Produced Image Highlights; Noisy Language, No $X$ ; Van Gogh . . . . . | 117 |
| 8.9  | Noisy Language, No $X$ Summary Charts & Examples . . . . .             | 118 |
| 8.10 | Noisy Language, No $X$ Summary Examples . . . . .                      | 119 |
| 8.11 | Colour Target Set . . . . .  | 123 |
| 8.12 | Experiment C1 Summary Charts & Examples . . . . .                      | 125 |
| 8.12 | Experiment C1 Summary Charts & Examples . . . . .                      | 126 |
| 8.13 | Experiment C2 Summary Charts & Examples . . . . .                      | 128 |
| 8.13 | Experiment C2 Summary Charts & Examples . . . . .                      | 129 |
| 8.14 | Experiment C2 Summary Charts & Examples; Noise Language . . . . .      | 131 |

|      |  |     |
|------|--|-----|
| 8.14 | Experiment C2 Summary Charts & Examples; Noise Language . . . . .  | 132 |
| 8.15 | Experiment C2 Summary Charts & Examples; Noise, No $X$ . . . . .   | 133 |
| 8.15 | Experiment C2 Summary Charts & Examples; Noise, No $X$ . . . . .   | 134 |
| 8.16 | Experiment C2 Highlights . . . . .                                 | 136 |
| 8.17 | Experiment C3 Summary Charts & Examples . . . . .                  | 137 |
| 8.17 | Experiment C3 Summary Charts & Examples . . . . .                  | 138 |
| 8.18 | Experiment C4 Summary Charts & Examples . . . . .                  | 141 |
| 8.18 | Experiment C4 Summary Charts & Examples . . . . .                  | 142 |
| 8.19 | Experiment C4 Summary Charts & Examples; Noise Language . . . . .  | 143 |
| 8.19 | Experiment C4 Summary Charts & Examples; Noise Language . . . . .  | 144 |
| 8.20 | Experiment C4 Grayscale Compositional Summary Charts & Examples    | 145 |
| 8.21 | Calculation of Spatial Frequency . . . . .                         | 146 |
| 8.22 | Debbie Ayles - Jesmond Barn . . . . .                              | 147 |
| 8.23 | Evolved Images With Uncomfortable Spatial Properties . . . . .     | 148 |
| 8.24 | Angular Spatial Frequency Analysis - Distance Variations . . . . . | 150 |
| 8.25 | Angular Spatial Frequency Analysis: Various . . . . .              | 151 |
| 8.26 | Angular Spatial Frequency Analysis: Jesmond Barn . . . . .         | 152 |
| A.1  | Target Image FFT Decompositions & Reconstructions . . . . .        | 168 |
| B.1  | Experiment J3 Compositional Summary Charts & Examples; $K = 10$    | 185 |
| B.1  | Experiment J3 Compositional Summary Charts & Examples; $K = 10$    | 186 |
| B.1  | Experiment J3 Compositional Summary Charts & Examples; $K = 10$    | 187 |
| B.2  | Experiment K3 Compositional Summary Charts & Examples; $K = 10$    | 188 |
| B.2  | Experiment K3 Compositional Summary Charts & Examples; $K = 10$    | 189 |
| B.2  | Experiment K3 Compositional Summary Charts & Examples; $K = 10$    | 190 |
| B.3  | Experiment P1 Compositional Summary Charts & Examples; $K = 10$    | 191 |
| B.3  | Experiment P1 Compositional Summary Charts & Examples; $K = 10$    | 192 |
| B.3  | Experiment P1 Compositional Summary Charts & Examples; $K = 10$    | 193 |
| B.4  | Experiment P1 Compositional Summary Charts & Examples; $K = 25$    | 194 |
| B.4  | Experiment P1 Compositional Summary Charts & Examples; $K = 25$    | 195 |
| B.4  | Experiment P1 Compositional Summary Charts & Examples; $K = 25$    | 196 |
| B.5  | Polar Coordinate Language Summary Charts & Examples; $K = 10$ .    | 197 |
| B.5  | Polar Coordinate Language Summary Charts & Examples; $K = 10$ .    | 198 |
| B.5  | Polar Coordinate Language Summary Charts & Examples; $K = 10$ .    | 199 |
| B.6  | Polar Coordinate Language Summary Charts & Examples; $K = 25$ .    | 200 |
| B.6  | Polar Coordinate Language Summary Charts & Examples; $K = 25$ .    | 201 |

|      |   |       |
|------|---|-------|
| B.6  | Polar Coordinate Language Summary Charts & Examples; $K = 25$     | . 202 |
| B.7  | Polar Coordinate Language Summary Charts & Examples; $K = 50$     | . 203 |
| B.7  | Polar Coordinate Language Summary Charts & Examples; $K = 50$     | . 204 |
| B.7  | Polar Coordinate Language Summary Charts & Examples; $K = 50$     | . 205 |
| B.8  | Circle, Grid, Offset Language Summary Charts & Examples; $K = 50$ | 207   |
| B.9  | Evolved Candidate Image With GP Tree - Grayscale . . . . .        | 208   |
| B.10 | Evolved Candidate Image With GP Tree - Colour . . . . .           | 209   |

# Chapter 1

## Introduction

Digital art can bring to mind many wide and varying concepts and examples. Many digitally-produced, original pieces have found their own acclaim, as have novel variations on existing classics. It is trivial for software to precisely replicate a digital image. Conversely, we find it difficult to autonomously produce new images with similar aspects. Forming correct abstractions between the digital data and visual interpretations is an ongoing challenge covering many fields of study.

Where a function or procedure is able to produce an image in a principled fashion, we can call these results procedural (procedurally generated) images [1]. The ability to form minor alterations in these procedures allows us to easily make changes in a structured manner, though it may not always be clear *a priori* how these changes will come to manifest. By combining together parts between the better performing procedural textures, we may gradually refine them, and allow them to exceed beyond the performance of any single prior texture. With this process of evolutionary refinement, we are able to explore many similar images which can feature novel and creative variation.

Procedural texture evolution has shown its use in applications ranging from interactive art systems [2], image filter adaptation [3], camouflage generation [4], and game asset generation [5] amongst others. Evolutionary algorithms (EA), and notably genetic programming (GP), are able to non-exhaustively explore the space of possible textures with little explicit understanding of how to effect high-level image changes [6][7]. Perhaps the most critical component in all EAs is the fitness measure, defining the meta-heuristic which guides the search to optimal solutions. With texture 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.

A number of measures for colour modelling have been explored in existing research, for both retrieval and synthesis, with various degrees of reported success [8][9][10]. Finding a measure for the comparison of spatial characteristics has remained difficult. We find that many existing frequency decomposition measures are not well-suited for the purpose of texture synthesis. It would therein be of great use to find a simple measure capable of guiding the shape of evolved images.

In investigating the existing measures that can be computed from a rendered texture, 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 in image analysis/retrieval research [11][12][13] have been found which use this measure and its aggregates to more effectively classify images based on computationally tricky but perceptively obvious attributes (*i.e.* Eastern vs. Western art; Portrait vs. Sketch vs. Landscape).

Between the successes found in texture synthesis with evolutionary algorithms, and the classification and perceptual relations with power spectral density and other Fourier decomposition abstractions, we believe the combination of these ideas may lead to improved texture generation schemes which can better target the high-level image characteristics of shape, texture, or style.

## 1.1 Purpose

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, we believe it shows great promise as a tool for guiding evolutionary textures with a more direct interpretation. Our goal is to explore the ability of these measures as used in evolutionary texture synthesis and evaluate the utility that these measures may have for the production of digital evolutionary art.

We will be considering our models of shape from a target image for use as a guide when evolving new images and textures. However, an exact reconstruction of a target image would not be ideal. Rather, 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.

In attempting to fulfil our goals, there are a number of milestones that will guide

our research. Largely as part of our background materials in Chapter 2.3, we aim to first refine understanding and intuitions regarding the 2D Fourier decomposition and its related, abstracted measures. Using our genetic programming engine with a basic configuration, we will then produce grayscale textures and explore the ability of various Fourier-based fitness measures to replicate spatial properties of target images. In this regard, these spatial measures should be considered at both the micro and macro scales. We will then evaluate the tractability of these measures for guiding both micro-level image refinements (style), and also lower frequency, macro-level image refinements (composition and shape). We hope that this will add an additional tool for these systems which can capture and guide these shape characteristics.

To facilitate a more objective evaluation of these measures, we will host and analyse an online questionnaire. The results of this questionnaire should show whether or not these measures can produce visually distinct styles or compositions which relate to their targets; measures which can therefore capture and guide critical spatial attributes in their evolved images. The questionnaire may be able to validate a link between our measures of spatial properties with human perception through the opinions of the participants. It is our hope that we will find efficient convergence toward these targets, and a means to generate textures with distinguishable spatial properties.

Finally, following any success we see while pursuing the above milestones, we will re-evaluate the utility of these measures for image composition in colour images, and using a more creatively-tuned evolutionary art system. We will display our exploration of power spectral density measures and their effectiveness when used for creating digital evolutionary art.

## 1.2 Thesis Structure

The layout of this thesis is roughly chronological and follows the flow of our exploration. In Chapter 2, we will present a summary of background material pertaining to genetic algorithms, procedural textures, and applications of the Fourier transform. Chapter 3 will then outline key pieces of literature which have guided or offered points of consideration for our experimentation. Chapter 4 will describe our system and its base configuration details, which will be followed by the initial fitness experiments in Chapter 5. Chapter 7 will break to the user validation survey —confirming the findings of the previous chapter —before returning to additional exploration and evaluations of utility for evolutionary artwork in Chapter 8. Potential concepts for

future studies will be outlined in Chapter 9.2, shortly after we conclude the results of our findings in Chapter 9.

# Chapter 2

## Background

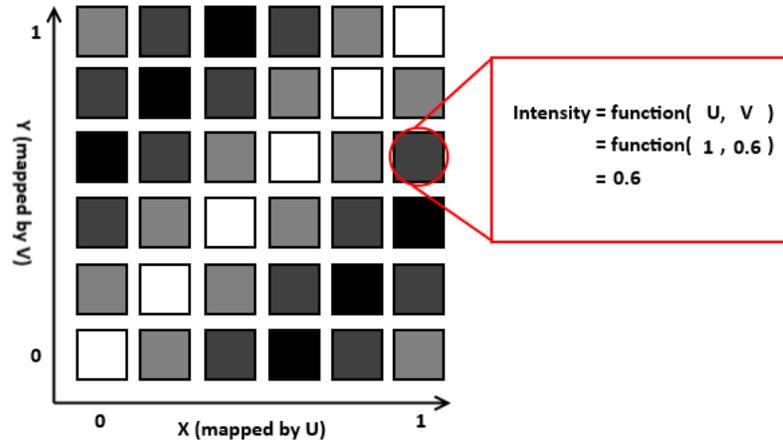
### 2.1 Procedural Textures

As the state of the art around computer graphics has advanced, we have seen greater and more novel methods of rendering scenes through the use of digital controls. Algorithmic methods to synthesize or alter various parts of a digital scene — including geometry and models, scene composition, shading, texturing — constitute the core premise of procedural rendering techniques.

While many aspects of digital rendering may have an algorithmic component to them, we consider textures and images to be “procedural” when they can be synthesized without direct reference to some base image [1]. In contrast to approaches like non-photorealistic filters [14], or other styled filters which apply effects on a source image, procedural textures should be able to produce a rendering without dependence on existing digital renderings.

Procedural textures have some potential advantages over using existing images in certain applications, such as compactness, variability through parametrization, and the ability to fill arbitrarily sized regions without obvious repetition [1]. While a potential advantage of having no fixed resolution is maintained, many of the other advantages will be lost for the exploration in this thesis. Rather, the key advantage of procedural textures for our usage is their efficiency and tractability to be adjusted with evolutionary algorithms, as discussed further in Subsection 2.2.1, and Section 3.1.

Figure 2.1: Rendering A Procedural Texture From A Function



### 2.1.1 Functional, Symbolic Expression

A procedural texture can use a formula or algorithm to determine colour or intensity values for each position in an image or texture (texel) [1]. While algorithms which produce textures of fixed size and shape are also common, we will concern ourselves most with functional textures which are rendered as a computation against the position corresponding to a texture element (texel) as in Equation 2.1. This function can be consecutively evaluated for each position in a digital canvas, determining its intensity or colour values, and thus producing the intended rendering (see Figure 2.1).

$$I = \text{TextureFormula}(U, V) \quad (2.1)$$

$$(R, G, B) = \text{TextureFormula}(U, V) \quad (2.2)$$

$$(R, G, B) = (r, g, b) \text{ where} \quad (2.3)$$

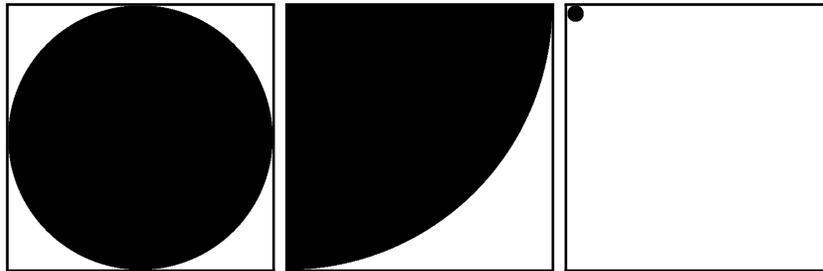
$$r = \text{TextureFormula}_R(U, V)$$

$$g = \text{TextureFormula}_G(U, V)$$

$$b = \text{TextureFormula}_B(U, V)$$

Simple 2D texture functions can be created using texel coordinate variables  $\{U, V\}$ , and a minimal set of arithmetic operators. These operators can be nested to induce greater image complexity, where periodic functions tend to provide image repetition, and discontinuous functions such as conditional operators can create contrasting intensity changes. Limiting the variable ranges, as well as operator and operand type can also influence the final rendering. For procedural textures, the texture coordinate

Figure 2.2: Rendering Window Bias — A unit circle rendered about the origin in windows, from left to right, of  $[-1, 1]$ ,  $[0, 1]$ ,  $[-1, 15]$ . Placement and scale show bias from window.



system, or rendering window, can refer to the range of numbers in each dimension which map between the integer indices of a pixel and the values provided as texels. Certain window ranges (e.g.  $[0,1]$ ,  $[-1,1]$ ,  $[0,128]$ , ...) may favour different end results in conjunction with the available operators, as shown in Figure 2.2.

### 2.1.2 Noise Generation

The introduction of seemingly random characteristics within an image is one of the most important tools in the creation of aesthetic procedural textures. While certain classes of textures may benefit greatly from the regularity that mathematical expressions can offer, the realm of aesthetic textures often appreciates a hint of randomness. Natural-looking images, and textures inspired from nature often see a combination of repeated patterns with some noisy variations [1]. Noise generator functions can easily introduce stochastic elements to the rendered texture.

A truly random noise implementation (*i.e.* providing a random value at each evaluation of the texture function) is often too chaotic to be used for aesthetic purposes. The contrast in values at adjacent positions is often too great. A more suitable result can be obtained by creating a smaller sampling of points to form a lattice, which can then be interpolated between [1]. This has an advantage of providing a highly stochastic feel, while maintaining a flowing, structured appearance — gradient noise. Such a noise generator should be seeded, and not truly random, since adjacent rendered positions are evaluated through separate calls to the procedural texture function. If the randomness is not reproducible, the procedural texture may evaluate with vastly different appearances in subsequent renderings.

Fractal noise is argued to be the most important element in modern, aesthetic texture generation [1]. Expanding upon a gradient noise implementation (such as the renowned Simplex or Perlin variants [15]), derivative noise generators can be produced

Figure 2.3: Fractal Noise Generation — FractalSum noise generated with 1 to 4 octaves

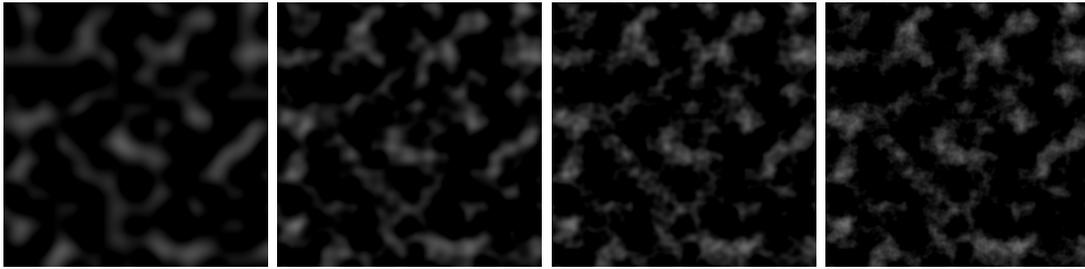
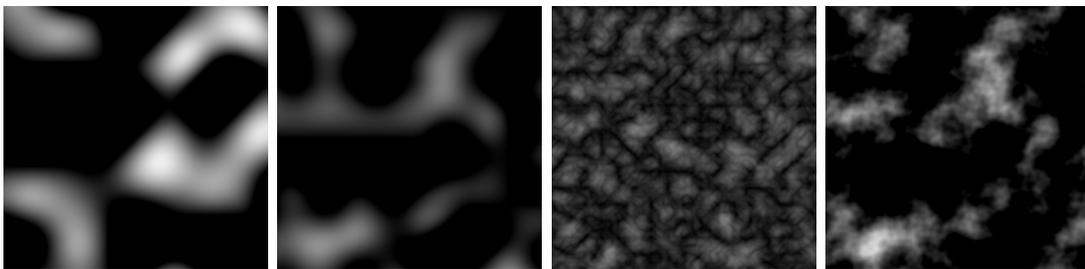


Figure 2.4: Noise Generator Examples — From left to right: Simplex, Perlin, Turbulence, Marble



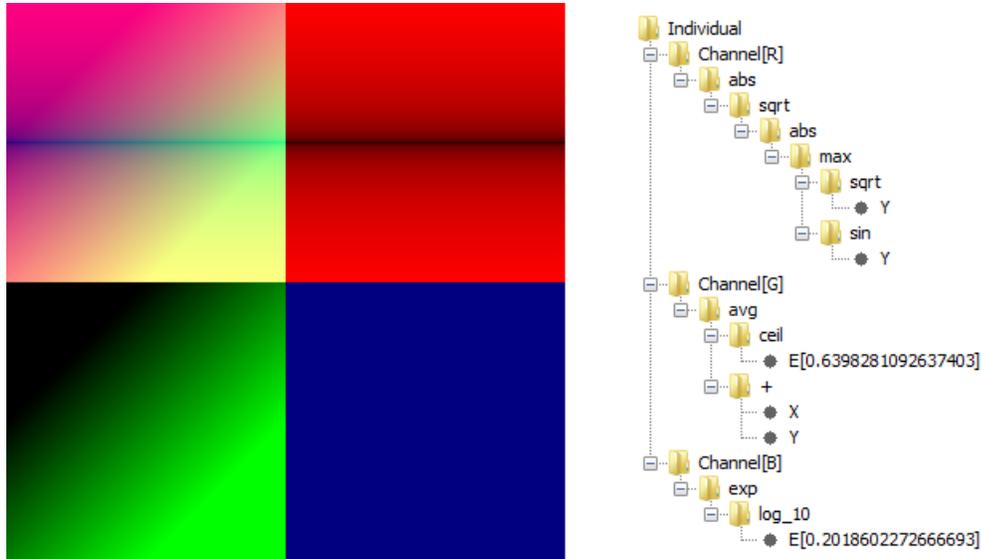
by combining resultant noise at various resolutions. By summing these intermediate noise results, we can generate a fractal noise that has both a soft flow and high amount of finer detail [16]. Figure 2.3 shows a smooth FractalSum noise generated at various maximum resolutions. Noise functions could be further adjusted and composed to alter the characteristics of the noise, achieving desired natural effects. Some examples are provided in Figure 2.4.

### 2.1.3 Colour Schemes

Colour is a key factor for the development of modern textures. Multiple colour channels can be evaluated separately, each with their own unique expressions (such as in Figure 2.5, and Equation 2.3). Alternately, an expression could use functions which support vectors, tuples, or other higher-dimensional types, and could consequently return a suitable tuple result for each position (*i.e.* a triplet for data in an RGB colour scheme, as in Equation 2.2).

In considering possibilities for colour schemes, the de-facto standard is RGB, which represents the intensities of red, green, and blue light to display for a given pixel. While the scheme is convenient for modelling the additive colour properties of projected lights, it can fall short in comparison to other schemes for modelling human

Figure 2.5: An example of a procedural texture produced through separate colour channel evaluations. Each colour channel of the texture is rendered through evaluation of a separate program tree. The final texture (top-left quadrant) is shown alongside renderings of each individual channel.



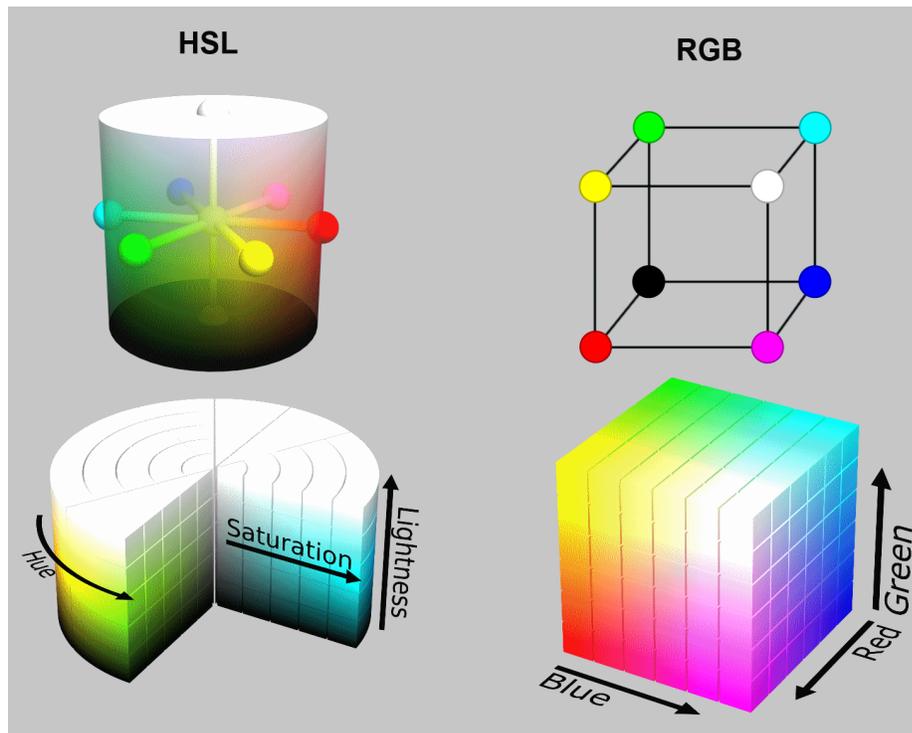
visual sensitivity, or for efficiently encoding the more design-focused properties of a colour palette. While RGB channels are simple for use at a basic hardware level, it would be expected that common gradients perceived across an image would induce non-trivial changes to each channel.

One of many alternative colour schemes for consideration is the HSL scheme [17], which includes a colour channel for each of the abstracted properties of hue, saturation, and lightness. Depending on the desired effect of the procedural texture, it may be easier to represent a desired effect by changing one of the more abstracted HSL channels, instead of requiring a change to each of the RGB channels in lock-step. Figure 2.6 shows a visual layout of the RGB and HSL colourspaces across their three dimensions. The translation from any colour scheme to the native RGB representation may require a minor cost in computation, but could be offset from the more conceptually friendly (and easier to manipulate) texture function.

## 2.2 Genetic Algorithms

Genetic algorithms (GA), and more specifically genetic programming (GP), are types of meta-heuristic, evolutionary algorithms which refine populations of possible candidate solutions to produce an improved approximate answer. In nature, we find

Figure 2.6: HSL and RGB Colour Models Compared[18][19]



the process of natural selection leads to generations of a species which have been passed more ideal genetic traits, and are better able to survive and adapt to their environments. By using this as inspiration, and continuously preferring the more fit individuals as candidate solutions, genetic algorithms are likewise able to gradually refine their population of final generated solutions over a number of virtual generations.

As genetic algorithms do not inherently make any assumptions about the data on which they operate, one key advantage is that we can use GAs to investigate a search-space of otherwise unknown structure. That is, GAs provide us a method to refine solutions where it may not be clear how to strictly improve upon an existing solution, so long as an evaluation criteria is provided. Genetic algorithms have shown their capability to solve problems of optimization, regression, classification, and artificial life among other search problems [20].

While an advantage to genetic algorithms is their utility in being purposed toward a large variety of problems, there are a number of elements requiring consideration and tuning to achieve useful results. As the algorithm refines the individuals in a population and improves approximation, it is necessary to determine a termination condition. It is possible that a single optimal solution can be found, at which point

we should not persist in searching further. However, it is not always the case that a perfect solution can be found, or even exist. Alternative termination criteria may be to stop when a “close enough” solution is found, after a number of generations has passed, or after some amount of time has elapsed.

Depending on the termination criteria, a maximum number of generations may need to be specified. A larger number of generations permits for more adjustments to a candidate solution, and thus a wider traversal of the problem search space, but we often see diminishing returns. The maximum number of generations would need to be balanced with the time it takes to execute the run, as better results may be found from performing multiple shorter executions.

A key configuration affecting performance and execution time of the algorithm would be the size of the population being refined. Reducing the number of possible candidate solutions to evaluate and refine each generation may reduce the execution time, but it likewise reduces the amount of genetic material that can be easily transferred amongst the remaining solutions. This is a problem of ensuring sufficient genetic diversity, and can be alleviated somewhat by keeping a copy of the best solutions of the previous generation (elitism), or by employing schemes with sub-populations, such as segregating and infrequently mixing groups of the individual solutions.

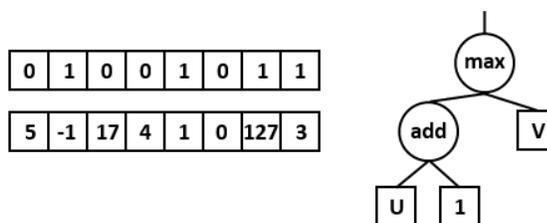
Other key considerations which we will explore in more depth include the representation of candidate solutions, possibility and frequency of solution-adjusting operations, and the evaluation criteria for how ideal or “fit” a solution is.

### 2.2.1 Representation

A critical requirement for the success of the genetic algorithms is to ensure that our representation of a possible solution be suitable for the type of problem we are solving. The way we encode an individual should be able to represent a valid solution to the problem, and should permit verification of the solution without error. For the sake of performance, it is ideal that a representation be able to hold and express easily mutable elements which alter small, granular aspects of the solution to allow refinement.

A canonical genetic algorithm representation is encoded in a flat array of some data type, where each index in the array may represent a small aspect of the solution. In the case of genetic programming, we consider the underlying representation of each individual to be one or more trees of functions and operands. Often for intermediate

Figure 2.7: Genetic Algorithm and Genetic Programming Representation – A flat GA representation using both binary (top-left) and signed byte (bottom-left) data types, and a GP tree representation (right).



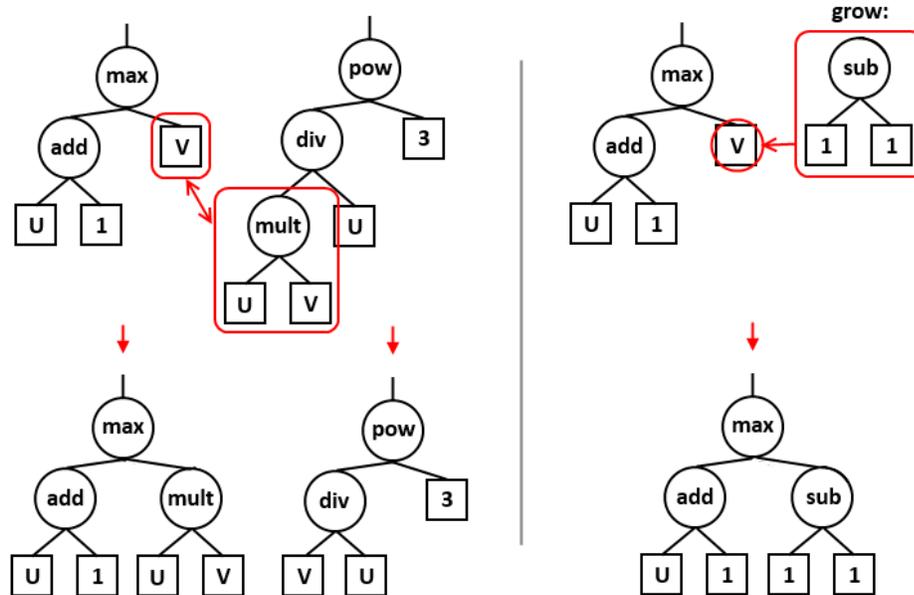
processing, we can view this representation as an S-expression. Figure 2.7 contrasts the canonical GA and GP representations.

An additional consideration for using a functional, tree representation pertains to language choices. The set of possible functions which can be used for a given node should be capable of expressing and leading candidate solutions to a valid answer (*e.g.* it would be difficult to express a periodic function using only basic arithmetic). Likewise, the set of possible terminal nodes can also bias solutions toward chosen coefficients, or may prevent a useful function from being formed (*e.g.* a classification problem missing a terminal which represents an input variable may not have sufficient ability to make a determination). Regarding terminals, we may often see input variables (and values derived from them), user-defined constants, and ephemeral random constants (ERCs, randomized values which persist for the duration of the run).

A concern in genetic programming with immediate consequences is the assurance that a solution can be evaluated without critical error. As nodes of a tree are recursively evaluated, we need to ensure that the range of each child node is the domain expected of that operand. Evaluation of a function node must be protected against errors including division by zero, and roots of negative numbers, requiring the creation of safe variants of any offending function.

In the case of textures and images, representation of individuals, and particularly its bias in artistic criteria must be considered. The use of bitmaps as representations is often too fine and granular; we often wish to use some abstraction to move ideas onto our canvas. As procedural textures can be generated through function composition, a GP representation seems well suited. However, we must be mindful of language choices, as it is often said to be one of the greatest factors in biasing the style of the resultant images. The meta-heuristic abilities of GP have proven themselves capable of overcoming the trickiness in meaningfully adapting these low-level expressions to have higher-level features and thus make it a reasonable search strategy for textures.

Figure 2.8: Genetic Programming Reproduction Operators – selected nodes being swapped in a crossover (left), and newly grown nodes replacing the selected sub-tree during mutation (right).



## 2.2.2 Reproduction Operators

The two key types of reproductive operators used to carry genetic material to subsequent generations are crossover and mutation. A canonical GA or GP representation has very straightforward methods of implementing these reproductive operators, though a non-standard representation may require custom implementations dependent on their representations and constraints. In all cases, it is ideal to plan the representation accordingly so that reproductive operators always result in valid solutions. A fragile representation may not be left in a valid state at the conclusion of a crossover or mutation operation, though it is possible (albeit additional work) to perform a correction on the individual after the operation completes.

The purpose of the crossover operation is to exchange genetic material, and therein aspects of a solution, between two parent candidate solutions. A pair of child candidates are formed through recombination by first copying each parent, and then exchanging a subset of their genetic information (values in their representation). In the ideal case, the aspects of a solution borrowed from the other parent would integrate well and produce one or more individuals which outperforms the original parents. This operation is the crux of genetic algorithms, as it is precisely what permits superior candidates to share beneficial traits, and thus have the population converge on improved solutions.

Crossover operations in canonical GAs are performed by first establishing regions along the flat, fixed-length vector representation, as denoted by a number of points. These points may be fixed, or chosen randomly for every set of parents [6]. The contents of alternating regions are then exchanged to produce the child candidates for the next generation. With a GP tree representation, two non-root nodes are chosen at random (possibly constrained by data type, depth or size), and swapped. Barring additional constraints, two child candidates should be produced by recombination for each execution of a crossover operation.

The other common reproduction operation is mutation, the desired result of which is to introduce genetic diversity. Without an element of mutation, we would see populations quickly converge to show a small subset of the initial genetic information available [6]. By randomly changing small aspects of a solution, we reintroduce individuals into the population which are (we hope) at least not substantially worse than the average, but which employ a potentially unseen trait in their solution. By slowing the rate of convergence and maintaining genetic diversity, we dissuade the search from getting too focused on local optima.

Mutation operations in the canonical GA representation is quite simple: a random region of the vector is overwritten with random data. While marginally more complicated with GP, a random non-root node is selected (again, possibly constrained by size or depth), and a tree-building method is used to produce a replacement sub-tree whose size is similar to the one being replaced. With GP, further mutation variants and configurations are commonly employed, such as adjusting the probability of selecting a terminal or non-terminal node for replacement, or using specific percentage adjustments to modify the value of ephemeral constants.

A final consideration in the use of various reproductive operations is to determine the probabilities of executing each possible operation. An excess of crossover operation calls per generation may result in quicker, premature convergence on a sub-optimal solution. Conversely, using too high of a mutation probability may not allow for more ideal solutions to recombine, and may be similar to a random search or hill-climber approach.

### 2.2.3 Evaluation

As the core concept of genetic algorithms is to preserve the genetics or aspects of more ideal solutions in subsequent generations, it is critical that we have a way of determining which of multiple candidates are most ideal. Darwinian evolution is

principled in the notion of small, inherited changes improving the survivability of a species. To replicate the concept of survival, we need to induce competition between individuals based on a function evaluating their “fitness”. These fitness functions are used to assign a score of correctness to each candidate solution, and must be crafted carefully.

The immediate use for the fitness function within the algorithm is for the selection of individuals whose genetic material is to be (partially) used for the next generation. Multiple strategies exist for choosing which parent individuals are used in a reproductive operation, but perhaps the most common method is a tournament selection scheme. In such a scheme, a fixed number of random individuals within the population are considered, and the one whose fitness function produces the best evaluation undergoes reproduction to the next generation. The number of individuals chosen for the tournament can in turn adjust the selection pressure, and so should be balanced between favouring the best candidates of the generation and maintaining genetic diversity.

In supervised learning systems, we may wish to evaluate a solution across numerous scenarios, positions, or data points to find an aggregate performance measure of the candidate solution. Often in these cases, we would like to withhold a portion of the data points for validation and testing; these points ensure that the solution is not over-trained, and do not contribute to a solution being carried over (or not) to the next generation. In unsupervised learning, such as in an art system, we are often less concerned with over-training and adaptability of our solutions, but instead hope for an exploratory search where fitness models guide solutions to many possible solutions in the search space having certain traits.

While supervised and unsupervised learning problems alike can make use of fitness functions which are able to automatically score candidate solutions, there is a notable alternative approach often seen when evaluating textures, images, sounds, and other subjective mediums, namely to have candidates evaluated with one or more human observations in what is known as aesthetic fitness scoring. While notable efforts have certainly been made to produce models of aesthetics, it is still very difficult to find precise definitions and quantitative scores for any sort of aesthetic criteria. By allowing users to specify their preference for each individual, we can further guide an exploratory search to match the interests of the user. A critical limit in the utility of this approach is that numerous subsequent evaluations can leave a user fatigued and drain their interest [21, 22, 23].

## 2.2.4 Multi-Objective Schemes

While some problems permit us to evaluate solutions with a simple measurement (such as overall error in a regression problem), there are also problems where a score cannot be easily expressed in a simple scalar value. Often in an exploratory search, we may want to find solutions which fulfil multiple constraints. When we wish to explore a problem landscape using more than one criteria, genetic algorithms require us to employ a multi-objective fitness scheme to express which individuals are more desirable during a comparison.

Simple schemes for combining multiple objectives into a single measure include weighted sum, and normalized product approaches. To produce a weighted sum, we simply consider providing a weight coefficient to each measure which determines how much that objective contributes to the final score, and summing these weighted values together. For a normalized product approach, individual measures in a  $[0, 1]$  range are multiplied together to form a compounded product. This would ensure that a poor score in any individual measure reduces the final resultant score appropriately. One considerable downfall to both of these simple schemes is that having non-normalized individual measures makes the schemes either unsuitable or heavily biased. While a weighted sum approach can operate with non-normalized data, it would then also require either all measures to be increasing with performance, or all measures to be decreasing with performance. The precise weighting amongst objectives is also an extra consideration.

One well-noted scheme which overcomes many of the issues in the previously discussed naïve approaches is the Pareto ranking [24]. Instead of using a single concrete evaluation to determine the fitness of the individual for the entire run, we instead score the individual in relation to the other individuals considered in that generation. We only consider an individual to be better than another if it is at least as good for every considered objective, and also better than the other at one or more of the objectives. This approach uses the concept of Pareto dominance to produce a number of tiers, or “fronts”, or equally viable solutions. Possible concerns with this approach are that it is not suitable for large number of objectives (frequently limited to 4), and that it can lead to solutions which excel in single objectives at the detriment of others.

A scheme which employs the best aspects of the above approaches is a sum of ranks (or average rank)[25][26]. After obtaining the raw measures for each fitness objective, we can produce a comparison of individuals by checking how each measure ranks amongst other individuals in the generation. By normalizing the rank of each

measure, we can then revert to a (weighted) sum of these rankings. This eliminates our requirement of normalized fitness measures, and the requirement of all objectives being strictly minimizing or maximizing, while also encouraging solutions which perform well across all considered objectives. Though this scheme tends to perform well with uniform weighting, we are also provided the option to customize this weighting to in turn reflect the importance of specific objectives. This approach has also seen improved results for texture generation when combining simple design criteria and aesthetics [27][28].

## 2.3 Fourier Transform

Fourier analysis is a well-known tool which sees substantial use in signal processing applications [29]. The operation 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 discrete Fourier transform (DFT) variant, and specifically the fast Fourier transform (FFT) implementation, provide for an efficient conversion of short sampled signals to an overview of their frequency composition.

The Fourier transform provides us a way to display the signal as a sum of cosine and sine terms, where the frequency of each periodic term relates to a component frequency found in the signal (this is well-illustrated in Figure 2.9). We also have the overall signal amplitude adjustment provided for the  $0^{th}$  coefficient. The result of such a decomposition is typically encoded as a complex number for each frequency. This complex number (Equation 2.12) can be translated in conjunction with Euler's formula (Equation 2.7) to provide the individual amplitudes of each sine and cosine term at that frequency. We will follow [31] for a brief outline of the derivation.

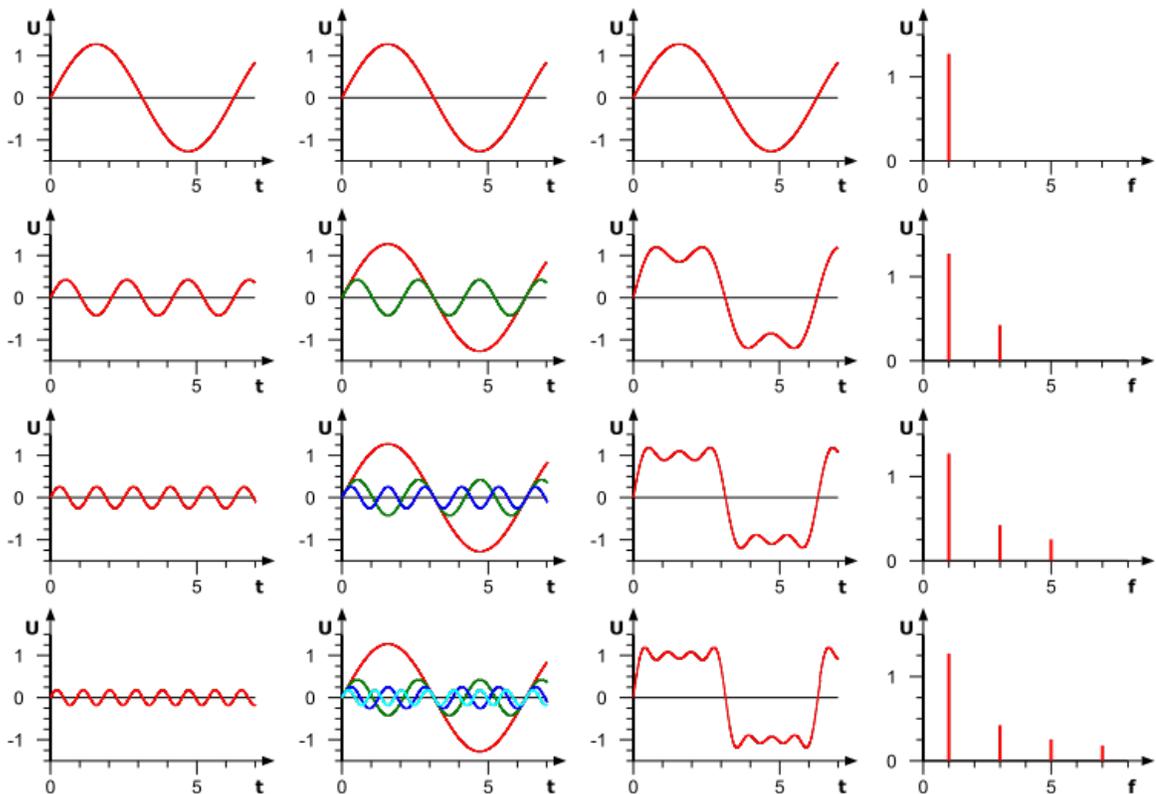
$$\omega = \frac{2\pi}{T} \tag{2.4}$$

$$\begin{aligned} f(t) = & a_0 + a_1 \cos(1\omega t) + b_1 \sin(1\omega t) \\ & + a_2 \cos(2\omega t) + b_2 \sin(2\omega t) \\ & + \dots \end{aligned} \tag{2.5}$$

$$\begin{aligned} & + a_n \cos(n\omega t) + b_n \sin(n\omega t) \\ = & a_0 + \sum_{n=1}^{\infty} (a_n \cos(n\omega t) + b_n \sin(n\omega t)) \end{aligned} \tag{2.6}$$

Figure 2.9: Fourier Synthesis in 1D: [30]

Each row introduces additional frequencies being added into the cumulative signal. We can see how any signal can be represented by the sum of component frequencies. From left to right, columns present: a pure frequency tone, cumulative signals overlaid, cumulative signals summed, Fourier analysis of component frequencies. The first three columns show signal amplitude  $u$  plotted over time  $t$ , with a base period (and signal length) of  $2\pi$  seconds. The last column shows the mapped amplitude  $u$  over the component cosines of frequencies  $f$ . We can observe that the first pure tone which has exactly one cycle (within in our measured signal length) shows high amplitude in its Fourier analysis at a frequency of 1. Likewise, the addition of other pure tones which repeat  $n$  times show some amplitude in their Fourier analyses at a frequency of  $n$ .



We can see specific cosine and sine amplitudes being summed to form a single signal in Equation 2.5, expressed succinctly in Equation 2.6. The complete signal function  $f(t)$ , is produced as a function over time or spatial samples, by summing all its pure component tones. Each tone of frequency  $n$  is scaled by their respective amplitudes  $a_n$  and  $b_n$  (for cosine and sine contributions), where the measure of angular velocity  $\omega$  (Equation 2.4) ensures the sine and cosine oscillations are measured relative to the observed signal length, or fundamental period  $T$ .

$$e^{in\omega t} = \cos(n\omega t) + i \sin(n\omega t) \quad (2.7)$$

$$\begin{aligned} f(t) &= a_0 + \sum_{n=1}^{\infty} \left( \frac{a_n}{2}(e^{in\omega t} + e^{-in\omega t}) + \frac{b_n}{2i}(e^{in\omega t} - e^{-in\omega t}) \right) \\ &= a_0 + \sum_{n=1}^{\infty} \left( \frac{1}{2}(a_n - ib_n)e^{in\omega t} + \frac{1}{2}(a_n + ib_n)e^{-in\omega t} \right) \\ &= a_0 + \sum_{n=1}^{\infty} (A_n e^{in\omega t} + B_n e^{-in\omega t}) \end{aligned} \quad (2.8)$$

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

$$B_n = \frac{1}{2}(a_n + ib_n)$$

With the use of Euler's formula (Equation 2.7), we are able to redefine coefficient terms into the complex plane, removing periodic functions from our definition (Equation 2.8).

$$\begin{aligned} B_{(n)} &= \frac{1}{T} \int_0^T f(t) e^{i(-n)\omega t} dt \\ &= \frac{1}{T} \int_0^T f(t) e^{-in\omega t} dt \\ &= A_n \end{aligned} \quad (2.9)$$

Observing the symmetry in Equation 2.9), we see that  $f(t)$  can be described through a single series of complex terms  $A_n$ . To mark the inclusion of  $a_0$ , and the broadened range of the series, we express this new series of complex terms as  $C_n$ . We are now able to adjust the range of our sum, and produce the elegantly succinct formula (Equation 2.10) which defines the Fourier series in complex terms (Equation 2.12).

$$f(t) = \sum_{n=-\infty}^{\infty} C_n e^{in\omega t} \quad (2.10)$$

$$C_n = \frac{1}{T} \int_0^T f(t) e^{-in\omega t} dt \quad (2.11)$$

$$= \frac{1}{2}(a_n - ib_n) \quad (2.12)$$

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.

Interpreting this in a visual application may be to consider decomposing the signal into repeating gradients. This closely related interpretation is referred to as discrete cosine decomposition (DCT), although canonical DCT has purely real values, and so discards any capture of frequency phase.

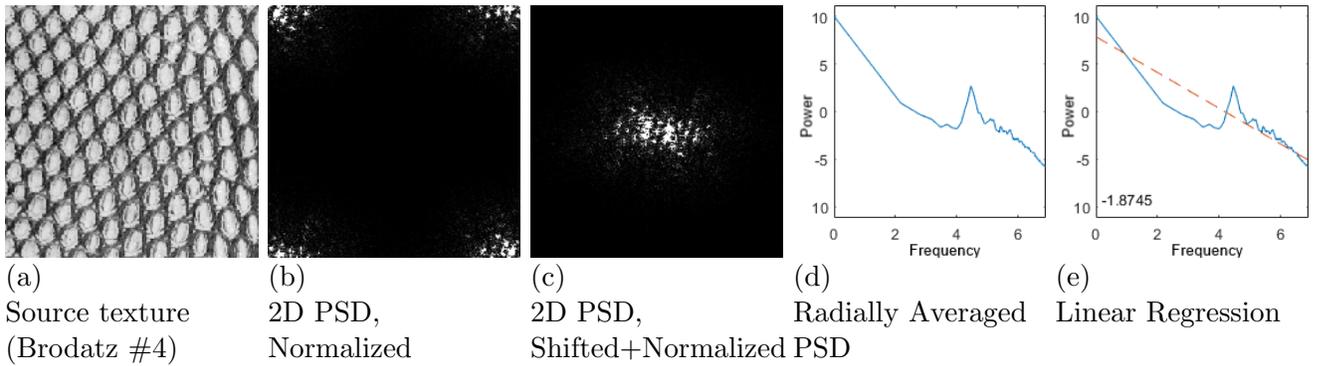
Adapting the Fourier transform to a 2D or higher-dimensional space can be found by applying the DFT on each index of the first dimension, and then again along each row of the results. This gives us the amplitude and phase of how each frequency contributes to the total 2D signal. In applications with images, we often see most of the high-energy coefficients appear around the central positions and main axes of the shifted FFT[32], as seen in Figure 2.10 and Figure 5.1. Where the amplitude of an audio signal may have an intuitive correspondence with sound wave pressure, amplitudes for a 2D image will be a measured in relation to their pixel intensity (typically 0..255), 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 [33]. In consideration of applying the FFT 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 [34] might assist in this matter.

One expectation held by the Fourier transform is that the samples it is operating on are from a periodic signal. If we were to apply the Fourier transform on a set of samples whose main frequency was not a product of our sample size, often identified from discontinuities between the first and last sample values, we would be violating this expectation. What might appear as artifacts (or, spectral leakage) would be

Figure 2.10: Power Spectra Pipeline

We begin by viewing the source image in subfigure (a), which undergoes 2D Fourier analysis, and 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).



all the unexpected frequencies needed to compensate for the disjoint produced while repeating the signal. Various functions known as windowing functions, or windows, can taper or otherwise adjust the shape of the signal ends to known values (often zero). To assist in reducing some of the smaller discontinuities in the signal, it is often advised to first filter it through an appropriate window function before continuing with further spectral analysis.

### 2.3.1 Power Spectral Density

The power spectral density (PSD) —or, power spectra —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 [35]. 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 Equation 2.13 below.

$$P_j = \left( \frac{C_j \overline{C_j}}{n^2} \right) \quad (2.13)$$

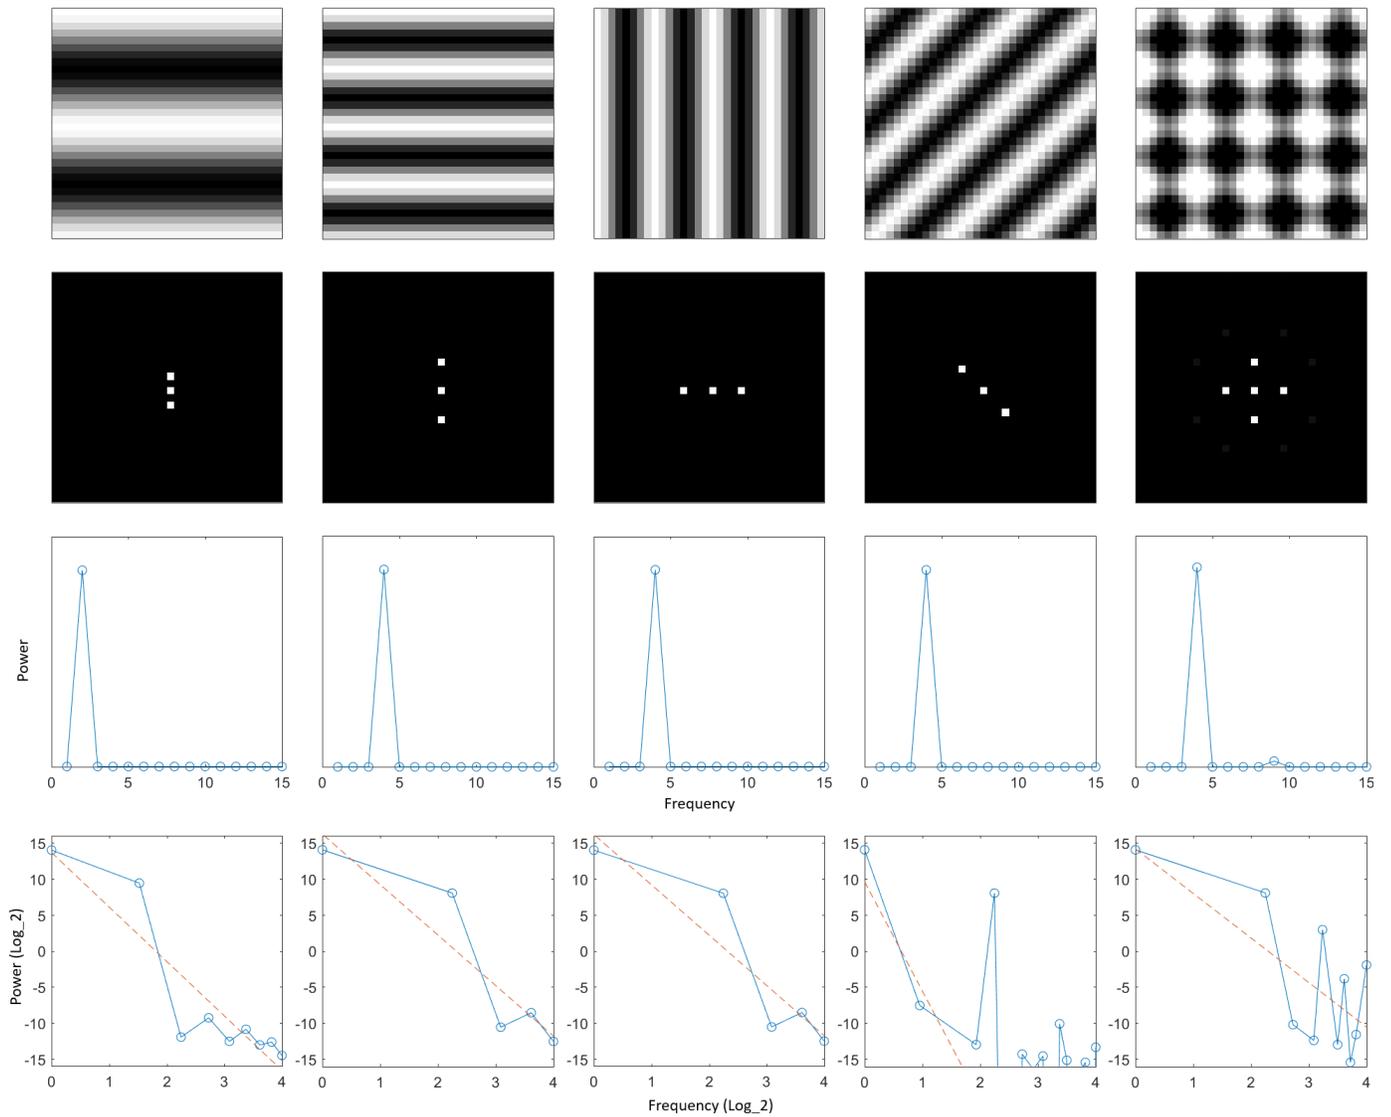
$$= \left( \frac{|C_j|}{n} \right)^2 \quad (2.14)$$

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 2.10. Between the DFT and the radial averaging methods, the power spectral estimate measure has the benefit of being approximately equal across rotation, and 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.* [12][13]) that “power spectral density” and related terms often refer to the radial average or similar abstractions.

Figure 2.11 assists in interpreting the 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) are contained at the center of the shifted FFT power coefficient display. Our first column shows a wave whose period is half of the canvas (input signal), and so the charted radially-averaged power spectra 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 artifacts 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 spectra 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. We strongly recommend those with continued interest refer to [36].

Figure 2.11: Power Spectra Interpretation &amp; Reconstruction

Columns 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



### 2.3.2 Compare & Contrast: Wavelets

Another common tool for signal analysis with many similarities to Fourier decomposition is wavelets [29]. While we do not particularly explore further use of wavelets, a number of pieces of background literature — some of which were highly inspirational for our work — have considered wavelets as an alternative means of producing spatial analysis. The process of wavelet analysis effectively reduces an input signal into component waves of various known duration – the eponymous “wavelets”.

Both the discrete Fourier transform (DFT) and discrete wavelet transforms (DWT) deconstruct a signal into sums of basis functions. In the case of Fourier, we have seen sine and cosine functions used, though we are likely to see more complex scale-varying basis functions (*i.e.* Harr, Meyer, Daubechies) used by wavelets. Efficient implementations can be found for Fourier and wavelet transforms with algorithmic complexities of  $O(n \cdot \log(n))$  and  $O(n)$  respectively. Both analyses aim to provide an overview of a signal’s component frequencies, but where Fourier analysis provides a more precise evaluation of the entire signal, wavelet analysis iteratively evaluates scaled subsections of the signal along set break points. This appears to be quite suited to capturing edge information and specific features.

Wavelets are often praised for their ability to be localized in both space and time [37]. While DFT can be robust to the scale of its input image, it is less resilient to changes which adjust individual pixels/frequencies. Many coefficients would need to be adjusted to accommodate a minor change to the source image, and vice-versa. Conversely, wavelet analysis performs decompositions iteratively at decreasing scale and fixed lengths. Changes to a position in an image would require adjustment to only spatially local coefficients (albeit at multiple resolutions). While some precision is lost in the frequency spectrum, the addition of having an approximate time-point more directly associable to the frequency coefficient is often seen as quite useful.

Wavelet analysis has found large support for applications centred about retrieval and classification, though we find that Fourier analysis is a more frequented measure in studies with visual similarity and perception. One possible reasoning for this may be found behind the works in *spatial frequency theory*, which suggest that human vision is decoded in a manner similar to Fourier analysis, measuring contrast frequencies along our field of view [38][39]. A number of articles have been found which consider the use of Fourier decomposition in their models of perception and relation to certain environments [40, 41, 42].

With respect to digital evolutionary art, wavelet analysis appears to be more frequented than Fourier analysis. The Genshade[43] and Gentropy[44] systems make

use of extracted wavelet coefficients to provide some spatial guidance in an unsupervised fitness approach. This technique was first outlined and borrowed from Jacobs *et al.* [45], and we will consider adaptations of this approach using Fourier coefficients. The curious lack of Fourier analysis in unsupervised evolutionary art and the success shown by the Jacobs *et al.* measure in the previous systems are strong motivations for the strategies we will be exploring.

# Chapter 3

## Literature Review

### 3.1 Evolutionary Textures

The use of evolutionary algorithms for texture synthesis had been pioneered by Sims [7], and has since been the topic of much previous and ongoing studies, adopting both interactive and unsupervised approaches [46][43][28][47].

In the founding work by Karl Sims[7], a system was presented which, through the guidance of a user, could gradually manipulate sets of graphical shaders to produce textures and animations of images fitting a desired aesthetic. While user interaction was required to provide the necessary aesthetic judgements for evolutionary selection, this evolution-guided approach to image refinement was a first in providing a computational exploration of artistic images. The system by Sims would come to constitute the first of many such evolutionary art systems.

An early attempt in the transition to unsupervised approaches came from Baluja *et al.*[46]. 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 also since been revisited with the inclusion of multiple abstracted image measures with some reported success [48].

A critical successor to Sims’ work was the Genshade system by Ibrahim[43]. While the synthesis capabilities of Genshade focused on the refinement of Renderman system shaders, advances were made to see results while evolving textures in an unsupervised manner. Various forms of image analysis 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 traits of the targeted image.

The Gentropy system by Wiens and Ross[44] expands upon the unsupervised approach of Genshade by providing additional image analysis measures. With concerns that the Genshade representation may present further bias to the evolved images, the system returns to a genetic programming approach using a simple language set. While we will shortly explore a number of variations seen in modern evolutionary art systems, we see that similar unsupervised, GP-based approaches have remained popular.

With the widened abilities of modern texture synthesis approaches, a number of applications have seen successful adaptations. While it is in the interest of this thesis to produce textures and filters for their own merit, we have also seen practical systems being developed in the areas of camouflage [4], and game asset [5] generation.

Alongside and expanding upon the evolution of textures, we have seen interest in creative and evolutionary systems across the domains of music[49], modelling[50], and architecture[51], amongst others[52].

## 3.2 Power Spectral Density

An initial investigation shows that a sizeable amount of research has been done related to the use of power spectra as a tool for classification and image retrieval. Despite this, we see little has been done related to the use of power spectra as a target for texture synthesis, or NPR filter evolution. It is precisely because of this that we will explore potential strategies to adapt this tool for use in evolutionary art.

### 3.2.1 Context: Vision

One key argument for the use of power spectral density to guide perceptual properties comes from spatial frequency theory. The theory purports that a human or animal visual cortex operates through coding signals in relation to spatial frequencies observed (in contrast to edge and line detection). Support for this comes from stronger neural response to sine wave gratings (of arbitrary frequency) than to imagery of images or bars. Numerous studies have and continue to explore this concept [38, 39, 40, 53, 54].

A notable article with high potential for practical use comes from identifying uncomfortable images through contrast and frequency analysis. Based on work observing the sensitivity ranges of luminance contrast, Fernandez and Wilkins had explored the undesirable effect of certain image contrast frequencies in human perception [41]. Power spectra of an image's luminance was investigated, and certain frequency octaves

were found to provide higher ratings of perceptual discomfort (within 2 octaves of 3 cycles/degree). Future studies could find inspiration in adapting this phenomenon to an explicit aesthetic measure, and analysis on some of our produced images appears to corroborate these findings.

### 3.2.2 Context: Classification & Retrieval

It was found by Neumann and Gegenfurtner that power spectra, in addition to other abstract information measures (including colour distribution and luminance) did permit for a marginal increase in retrieval performance [10]. Work by Graham *et al.*[12] also explores and surveys a number of methods making use of power spectra slope for art classification.

Some works by Balboa[42] and Millane[55] were found which correlate classes of natural images to the slope of their power spectra. While we expect that the slope of the power spectra alone is insufficient to generate natural images, there is a strong suggestion that power spectra may be used to perceptually differentiate attributes of complex textures.

A number of articles by Graham reinforce the idea that power spectra can be used to assist in classification of art pieces [56][12]. Graham had concluded that most pieces of art fall in the same spectral range of natural images. While power spectra measures alone are not sufficient to limit the potential art pieces to those that are more aesthetically pleasing, they were able to distinguish with some accuracy between art piece origin [11], and certain classes of images (*i.e.* human portraits, landscapes, sketches) [13]. Readers interested in a larger survey of how using power spectral density may relate to perception should be directed to references section of [13].

## 3.3 Spatial Measures

The need for measures permitting comparison of spatial properties tends to get resolved through one of two main concepts. We see that frequently taken approaches either extract key features (and their positions) from a source or target image, or perform 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, eccen-

tricity, axis orientation, and algebraic central moment information [57].

In some of the earlier unsupervised evolutionary art systems, particularly Genshade[43] and Gentropy[44], spatial features were compared via extracted coefficients from wavelet-based analysis measures. The results from Wiens [44] may be one of the more positive suggestions for using a power spectra fitness approach.

Shortly prior to the Genshade system by Ibrahim, a notable paper pertaining to image retrieval was published by Jacobs *et al.*[45]. In this paper, a quick algorithm was proposed capable of extracting the key coefficients from a 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 may also prove effective in guiding evolutionary systems.

However, this jump to a wavelet-based analysis, while ultimately showing some effectiveness, may glance over alternative offerings. One considered problem with a frequency analysis approach is in the inability to effectively handle images with multiple colour channels [33]. One potential solution to this was proposed through the use of the quaternion Fourier transform [34], which does not have an immediately adaptable equivalence with wavelet analysis.

A criticism common to all types of frequency analysis lies in the fact that a perfect solution would exactly replicate the target image [58]. In evolutionary art, we are often not trying to get a precisely identical image to our target. If our system fails to capture the precise image, this is actually preferred; we only want to capture key aspects of the image, and explore the landscape of possible solutions which are in some way similar. While we could adjust our fitness evaluations to prefer some amount of error, we find 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.

# Chapter 4

## System Design

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

### 4.1 Power Spectral Density Measures

For all the various experiments conceived for this study, a number of PSD-related measures would need to be efficiently calculated. Across the levels of abstraction, we would need to compute and obtain the 2D power coefficient matrices, the radially-averaged power spectral density, and its linear regressions.

An early decision in the development of the system was to make use of MATLAB[59] to assist with computation of power spectral density measures. Tooling within MATLAB allows us to generate native C code, which could then be used in the Java-based evolutionary system (described below, Section 4.2) through use of the Java Native Interface (or commonly, JNI) framework. This permitted for an expedient start to our exploration, and then allowed for a focus on the refinement — rather than development — of these measures. A trade-off to this integration approach was the limitation of a single evaluation thread, as the produced native code could not be effectively and safely run in a multi-threaded manner. The increase of wall-clock time for a single run was exchanged for multiple concurrent runs, and was not a substantial concern during our experimentation, though improvements could be made for future work.

To evaluate and produce the 2D power spectra (and phase) coefficients, and then find the radially-averaged power measure, we use the code outlined in Table 4.1 and

Table 4.1: MATLAB Code for 2D FFT Power and Phase

```

1  function [power,phase] = FFT2DPowerPhase(img)
2      [N,M] = size(img);
3      imgf = fftshift(fft2(img));
4
5      %power = (imgf.*conj(imgf))/(N*M * N*M);
6      power = (abs(imgf)/(N*M)).^2; % normalize
7      phase = angle(imgf);          % find angle in complex plane
8  end

```

Table 4.2: MATLAB Code for 2D Radial Average

```

1  function [ravg] = RadialAverage(ftArr)
2      [N,M] = size(ftArr);
3      assert(N == M);
4
5      L = floor(N/2)+1;
6      ravg = zeros(1,L);
7      rpts = zeros(1,L);
8
9      for y = 1:M
10         for x = 1:N
11             % consider polar coordinates, average across same rho
12             [~,r] = cart2pol( x-(N/2)-1, y-(M/2)-1 );
13             r = round(r)+1;
14             if (r > L); continue; end;
15             ravg(r) = ravg(r) + ftArr(x,y);
16             rpts(r) = rpts(r) + 1;
17         end
18     end
19     for l = 1:L
20         if (rpts == 0); continue; end;
21         ravg(l) = ravg(l) ./ rpts(l);
22     end
23 end

```

Table 4.2 respectively. A discussion on the computation of 2D Fourier decomposition and power is covered in Section 2.3. Section 2.3.1 also briefly discusses the radial average of this measure, though the core concept of its implementation is to bin and average the power coefficients based on their computed distance from centre (through the use of a polar coordinate system, see Table 4.2, Line 12).

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 background literature. Charting of PSD throughout this report will use  $\log_{10}$  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 expectedly 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.

A decision made in the early stages of research was to forgo any 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 non-regular signals, such as typical non-repeating images, to reduce heavy artefacts in the decomposition, with the specific window functions and parameters dependent on the expected signal. The absence of any discussion related to windowing in most of background research seems like a potential oversight, but the first step for our measures was to closely reproduce the published values. This lack of windowing trades the capture of information at the image boundaries for a more complex set of coefficients in the decomposition. Insofar as both target and candidate images use the same windowing (or lack thereof), comparing the measures should still hold, though GP may find the match less tractable.

While much of the existing published research has sparse details in the minutiae of their computation and regression of PSD measures, we have found our library to produce results closely matching the literature outlined in Section 3.2, and specifically those from Graham *et al.* [13].

## 4.2 Genetic Programming Engine

The evolutionary art system we will use to generate textures is a custom extension of the ECJ library for Java [60].

As we have showed our motivation for using evolutionary algorithms (and specifically GP) in section 2.1.1, our evolutionary system will employ a genetic programming tree representation to evolve symbolic expressions for procedural textures. Much of our early experimentation focuses on spatial attributes of an image, and so we find that in our initial exploration, grayscale textures are not only adequate, but are indeed preferable over the artistic colour texture renderings. To suitably represent this, our GP individuals need only a single tree used to evaluate luminosity or intensity. Later experimentation will expand to colour textures, and at that time, we will expand our individuals to hold 3 trees; one tree will be used for each colour channel in the RGB colour-space.

Initially, a population size of 500 was used as per [61], but diminishing genetic diversity was found by mid-run. Doubling the population size, in addition to the inclusion of diversity-penalties with the sum-of-ranks fitness scheme, was seen to assist in maintaining genetic diversity. Generation count was determined empirically from early runs, based on [44], with further adjustment from the considerations noted below. It was seen that 30 runs would provide a sufficient basis for statistical significance on comparisons across similar fitness measures, while balancing the wall-clock time required to complete the computations. Later experiments have seen a decrease in run count due to the higher computational complexity and time required to complete, such as seen with the introduction of noise operators and increased colour channels. At this point, often 9 runs were completed to show result diversity with less concern to statistical evaluations.

The wall-clock run times for the system configured for basic grayscale textures was 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 being evaluated concurrently, though the parallel nature of the system 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, with coloured noisy textures requiring an approximate average of 12 hours for completion of a run.

While our GP genotype representation will remain as `tree(s)`, we must make some considerations for their evaluations into images. While there has been some ongo-

ing exploration into typed GP systems, this research will for now constrain itself to operands and operators dependent on double floating point precision (`double`) values. Each tree should evaluate to a single `double`-type value, which could be clamped to  $[0, 1]$  before being scaled according to colour channel ( $[0, 255]$  for grayscale and RGB). To defer worries regarding cropping or padding, a rendering size of 128 by 128 pixels was used before sending the image through other evaluation measures, as square power-of-two images were efficient for Fourier decomposition. Larger images were frequently produced for external purposes, so a normalized texture coordinate system of  $[-1, 1]$  was used across both rendering dimensions. While other coordinate systems could be used, it was considered that a 2D space with equal lengths about the origin coordinate might better display the effects of phase on potentially symmetrical textures.

Working with GP and tree representations, we consider three main types of reproductive operations. Mutation and crossover operators are staples of GP, and so we have included them as possible operators with 20% and 70% probability respectively, keeping them in similar proportion to work by Heijer & Eiben [47], though slightly higher than others [44]. The introduction of ephemeral value mutation allows for the randomized constants to be slightly altered by 1%, allowing for finer, hopefully less destructive, adjustments to the rendered image. The ERC mutation operator has been included at a probability of 10%, and is responsible for a proportional decrease in likelihood to execute the crossover operator. So as to remove the possibility of losing or destroying the best found individual in a generation, we have allowed elitism for the single best individual of a generation to be retained unaltered in the subsequent generation. A table of other parameters, including reproductive operator settings, is provided in Table 4.3.

For consistency, all experimentation was performed using a normalized sum-of-ranks fitness scheme. Work in [61] has shown summed multi-objective ranks as a fitness measure can perform well in balancing targeted properties of evolutionary art without the dominance of individual objectives. When used in a single-objective problem, such as our early experiments, a sum-of-ranks fitness model will still preserve ranking information across that one objective, which would not provide biased influence during tournament selection. A benefit to using sum-of-ranks on these problems is that we can maintain a consistent diversity penalty scheme across all experiments. For individuals whose ranks in all objectives are identical, the second individual will have a penalty of 10 added to each of their ranks. Additional individuals found with the same scores as the first will incrementally receive an additional penalty of 10 rank

Table 4.3: Genetic Programming Engine Parameters Overview

| Parameter                    | Value              |
|------------------------------|--------------------|
| General                      |                    |
| 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 |
| New Node Depth               | [2, 6]             |
| Grow Probability             | 50%                |
| Parent Selection             |                    |
| Selection Method             | Tournament         |
| Tournament Size              | 3                  |
| Node Selection               |                    |
| Terminals                    | 10%                |
| Non-Terminals                | 10%                |
| Reproductive Operators       |                    |
| Crossover                    | 70%                |
| Mutation                     | 20%                |
| ERC Mutation                 | 10%                |
| Settings: Crossover          |                    |
| Max Depth                    | 17                 |
| Attempts                     | 1                  |
| Settings: Mutation           |                    |
| Max Depth                    | 17                 |
| Attempts                     | 1                  |
| Builder                      | Grow               |
| New Node Depth               | [5, 5]             |

points (the fourth common individual would receive a total of +30, and so on). The hope of such a penalty would be to maintain genetic diversity in the population by penalizing perfectly identical results. While the high precision of some of the fitness measures explored might lessen the opportunities for these penalties to be effective, we expect that the penalties should assist in reducing the effects of GP bloating.

The termination criteria for a run is to complete 100 generations. While “perfect” individuals have been produced for some simple compositional targets, this is otherwise a difficult problem, where finding such a “perfect” solution is not reasonably expected. Further, with later multi-objective experiments, an evaluation of raw scores instead of objective ranks is required to continue with this termination criteria. In some experiments, more aesthetically pleasing results were found prior to convergence. This suggests that an extended generation cap or termination by some convergence criteria would not always be beneficial. While some exploration was performed with longer generation caps, it was found early on that most objectives suitably converged by about the 100th generation.

### 4.2.1 Texture Languages

We begin with a basic, largely mathematical and symbolic texture language. Noise and more complex functions will be considered later after some experimentation, but it was not believed in the initial exploration that more “creative” functions be required to reach a solution. It was of interest to see how well our fitness measures could guide spatial attributes without the more stochastic (and computationally expensive) noise functions. We will later reconfirm the importance of language choices when we see substantial performance increases to certain targets with the addition of polar coordinate system variables. An overview of the GP language terminals and non-terminals can be found in Table 4.4, and the various extensions to the language can be seen in Table 4.5.

As part of the extended language, a number of noise operators are included. There have been various schemes to handle noise and their respective seeds across different published GP systems. In our system, noise operators are treated as an extension of ephemeral random constants; noise seeds are generated at a GP node level, and can be mutated or adjusted the same as any ephemeral constant (including through the ERC mutation operator). This approach offers less destructive changes in reproduction over seeding noise at the individual level, and increased genetic diversity when compared to employing run-level seeds. The approach we have taken may see benefits from higher

Table 4.4: Genetic Programming Engine Base Language Overview

| Category     | Arity                                      | Display | Description  |      |   |
|--------------|--|---------|--|------|---|
| Variables    | 0  | X       | Horizontal rendering coordinate                    |      |   |
|              |  | Y       | Vertical rendering coordinate                      |      |   |
| 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     | Periodic, trigonometric sine                       |      |   |
|              |  | cos     | Periodic, trigonometric cosine                     |      |   |
|              |  | tan     | Periodic, trigonometric tangent                    |      |   |
|              |  | sqrt    | Square root  |      |   |
|              |  | exp     | $e$ (Euler's number) raised to the operand         |      |   |
|              |  | pow2    | The operand raised to the power 2                  |      |   |
|              |  | pow3    | The operand raised to the power 3                  |      |   |
|              |  | log_E   | Natural log  |      |   |
|              |  | log_10  | Log, base 10                                       |      |   |
|              |  | 2       | 2  | +    | Addition  |
|              |  |         |  | -    | Subtraction   |
| *            | Multiplication                             |         |  |      |   |
| /            | Safe division; a zero divisor returns zero |         |  |      |   |
| max          | The greater of two operands                |         |  |      |   |
| min          | The lesser of two operands                 |         |  |      |   |
| avg          | The mean of two operands                   |         |  |      |   |
| pow          | $arg[0]$ raised to $arg[1]$                |         |  |      |   |
| 3            | 3  |         |  | lerp | Linear interpolation between $arg[0]$ and $arg[1]$ based on normalized (clamped to [0, 1]) $arg[2]$ |
|              |  |         |  |      |   |
| Conditionals | 4  | IfGT    | If $arg[0] > arg[1]$ then $arg[2]$ , else $arg[3]$ |      |   |

Table 4.5: Genetic Programming Engine Extended Language Overview

| Category  | Arity | Display                        | Description  |
|-----------|-------|--------------------------------|--|
| Variables | 0     | <b>Rho</b>                     | Polar coordinate; distance from $\{0, 0\}$   |
|           |       | <b>Phi</b>                     | Polar coordinate; angle about $\{0, 0\}$ to $X$ axis   |
| Spatial   | 1     | <b>Circle</b>                  | Gives 1.0 where $Rho \leq arg[0]$ , otherwise 0.0  |
|           | 3     | <b>Shift</b>                   | Evaluates $arg[0]$ as if the rendering position was shifted by $arg[1]$ horizontally, and $arg[2]$ vertically  |
|           |       | <b>Tile</b>                    | Evaluates $arg[0]$ as if the rendering position was scaled and offset to tile horizontally for $arg[1]$ windows, and vertically for $arg[2]$ windows |
| Noise     | 0     | <b>Simplex</b> <sup>†</sup>    | Simplex noise generator  |
|           |       | <b>Marble</b> <sup>†</sup>     | Marble noise (see [1])   |
|           | 1     | <b>FractalSum</b> <sup>†</sup> | FractalSum/Smooth noise  |
|           |       | <b>Turbulence</b> <sup>†</sup> | Turbulence noise   |

<sup>†</sup> 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.

mutation, but should also be able to preserve their characteristics in a crossover recombination.

Optimized Perlin and simplex noise generators have been borrowed from [62] and [63] respectively. While two base noise generators are available, it was considered unnecessary to include both, so only the simplex noise generator was included due to its more ideal window range. The fractalsum, turbulence, and marble noises have been based on the Perlin noise implementation as originally conceived. For these noise variants, coordinate scaling has been applied to ensure noise is applied across the  $[-1, 1]$  rendering window.

# Chapter 5

## Initial Exploration

Our exploration begins with the naïve approach of trying to reduce error with the basic FFT measures. Before moving on to more involved strategies, we evaluate the effectiveness of the abstract measures, and gradually move to the higher-dimensional FFT coefficients. Having found these lacking, we adapt an existing approach found to extract and compare a subset of key coefficients, and adjust how errors related to phase are incorporated. With some promising fitness measures explored, we will further refine these measures for use in recreating compositional attributes.

One can follow Table 5.1 for an overview of the discussed notable experiments, which relates the experiments to their discussion section, and provides a brief description. The basic GP system and parameters being used throughout the experiments (unless otherwise noted) is presented in Section 4.2.

Table 5.1: Overview of Experimental Variations

| Section | Label | Description  | Number Of Objectives |       |        |       |
|---------|-------|--|----------------------|-------|--------|-------|
|         |       |  | Power                | Phase | Colour | Total |
| 5.2.1   | E1    | PSD regression slope target  | 1                    | —     | —      | 1     |
| 5.2.1   | E2    | PSD regression slope & offset targets  | 2                    | —     | —      | 2     |
| 5.2.2   | E3    | PSD radial average curve fitting target, sum of squared error  | 1                    | —     | —      | 1     |
| 5.2.3   | E4    | PSD coefficient array target, sum of squared error   | 1                    | —     | —      | 1     |
| 5.3.1   | J1    | Significant coefficient extraction mismatch as per Jacobs <i>et al.</i> [44][45]   | 1                    | —     | —      | 1     |
| 5.3.1   | J2    | From J1, significant coefficients ranked, sum of rank mismatch   | 1                    | —     | —      | 1     |
| 5.3.2   | J3    | Coefficients extracted as per J1, up to half mark loss for phase   | 1                    | †     | —      | 1     |
| 5.3.3   | K1    | Power coefficient error and phase angle error normalized as single objective   | 1                    | †     | —      | 1     |
| 5.3.3   | K2    | Power coefficient error and phase angle error as multiple objectives   | 1                    | 1     | —      | 2     |
| 5.3.3   | K3    | As O1, with the “0th coefficient” as its own objective   | 2                    | 1     | —      | 3     |
| 6.3     | P1    | As J3, with squared phase error as a separate objective. Phase error scaled to truncated rank, and non-matched positions take full Pi (scaled) | 1                    | 1     | —      | 2     |
| 8.2.2   | C1    | P1 across 3 colour channels (RGB)  | 1                    | 1     | x3 ‡   | 6     |
| 8.2.2   | C2    | P1 across 4 colour channels (Y,RGB)  | 1                    | 1     | x4 ‡   | 8     |
| 8.2.2   | C3    | P1 across 3 colour channels (HSL)  | 1                    | 1     | x3 ‡   | 6     |
| 8.2.3   | C4    | P1, and also CHISTQ — histogram difference weighted by colour similarity   | 1                    | 1     | 1      | 3     |

† No separate objective for phase, but included as penalty to power objective.

‡ Colour objectives are extracted as power and phase measures for each individual channel in the given colour scheme.

## 5.1 Genre Images

In creating a pool of possible image targets to choose from, the notion of separation by genre was borrowed from Graham *et al.*[13]. In attempting to define visually distinct genres, the following classifications, while still very broad, were considered:

- portrait photos
- abstract paintings ★
- object photos ★
- texture patterns (regular & irregular) ★
- landscape photos
- geometric patterns
- abstract photos
- flag elements
- portrait paintings ★
- cartoon elements
- object paintings
- design elements
- landscape paintings
- compositional elements (simple, abstract) ★

Genres sampled for targets used in our experiments have been marked with a ★.

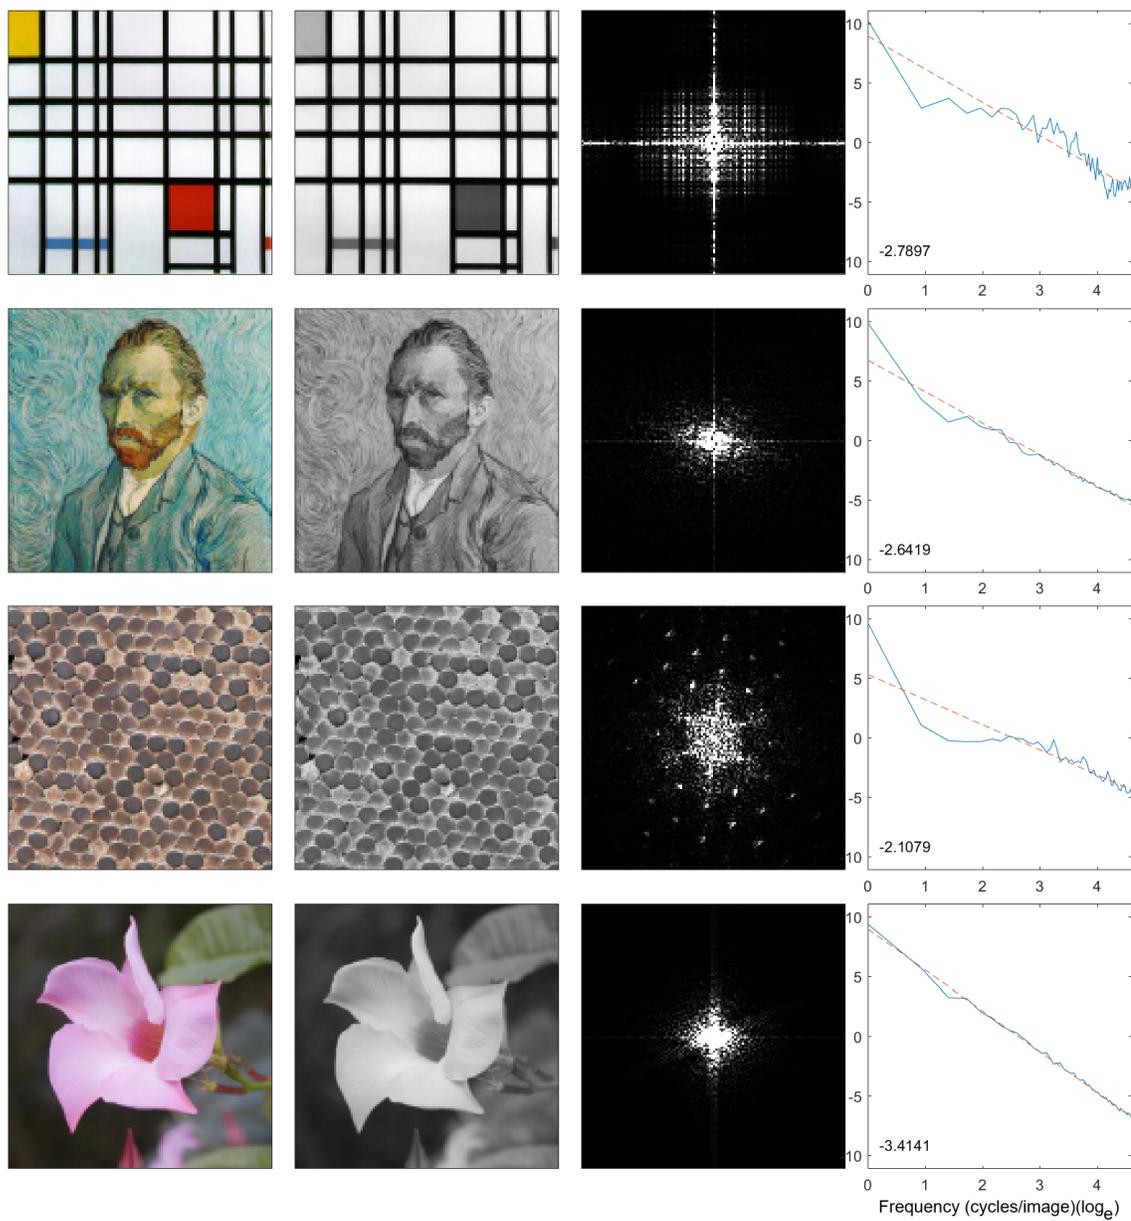
With the finite computing resources available, the initial selection of targets was limited to 4 visually distinct images across the genres. One abstract painting, one portrait painting, one texture pattern, and one object photo were chosen to balance the broad visual genres used with a reasonably workable number of images. With initial experimentation, targets were not chosen with concern to the difficulty a system would have in replicating similar visual features. Rather, traits desirable in target selection were the distinctness of an image amongst the other selected targets, the variety with regard to possible image genre, and different macro- and micro-level styling. Targets used for exploring compositional properties and integration with colour measures are explored later in Sections 6.1 and 8.2.1 respectively.

Initial targets are shown in Figure 5.1, and Figure A.1 in the appendix shows the target images with various Fourier analyses, and coefficient-limited reconstructions.

Figure 5.1: Initial target set. Featured top to bottom:

- Piet Mondrian – *Composition with Yellow, Red, and Blue*,
- Van Gogh – *Selbstporträt*,
- Procsilas Moscas – *Cable Ends - San Francisco Bay Bridge* [64],
- Vincent Gircys – an untitled photograph of a flower.

Target images are shown alongside their grayscale renditions, their power spectra coefficients, and radially-averaged power spectra charted with linear regression.



## 5.2 Simple Regression & Error

Our first four sets of experiments consider the error between FFT decompositions from evolved individuals and a target at various levels of abstraction. Beginning with the most abstract, and most prevalent in our literature review (*e.g.* [10][11][12][13][42][55]), we consider a fitness scheme which compares the difference of slope found through a linear regression. A simple expansion from there is to also incorporate the associated vertical offset measure produced from the regression. Gradually returning from abstractions to the raw measure, we consider fitness measures from squared error between radially averaged spectra, to squared error of FFT coefficient amplitudes.

To evaluate the targets and evolved candidate images, we can obtain measures at various levels of abstraction from our PSD library. The processing pipeline outlined by Figure 2.10 guides the measures we will use by virtue of being readily extracted for comparisons. Most of these initial fitness measures will use a single evolutionary objective when comparing between the target and each evolved candidate. Experiments in this section will make use of our standard, described evolutionary texture system with the basic GP language set.

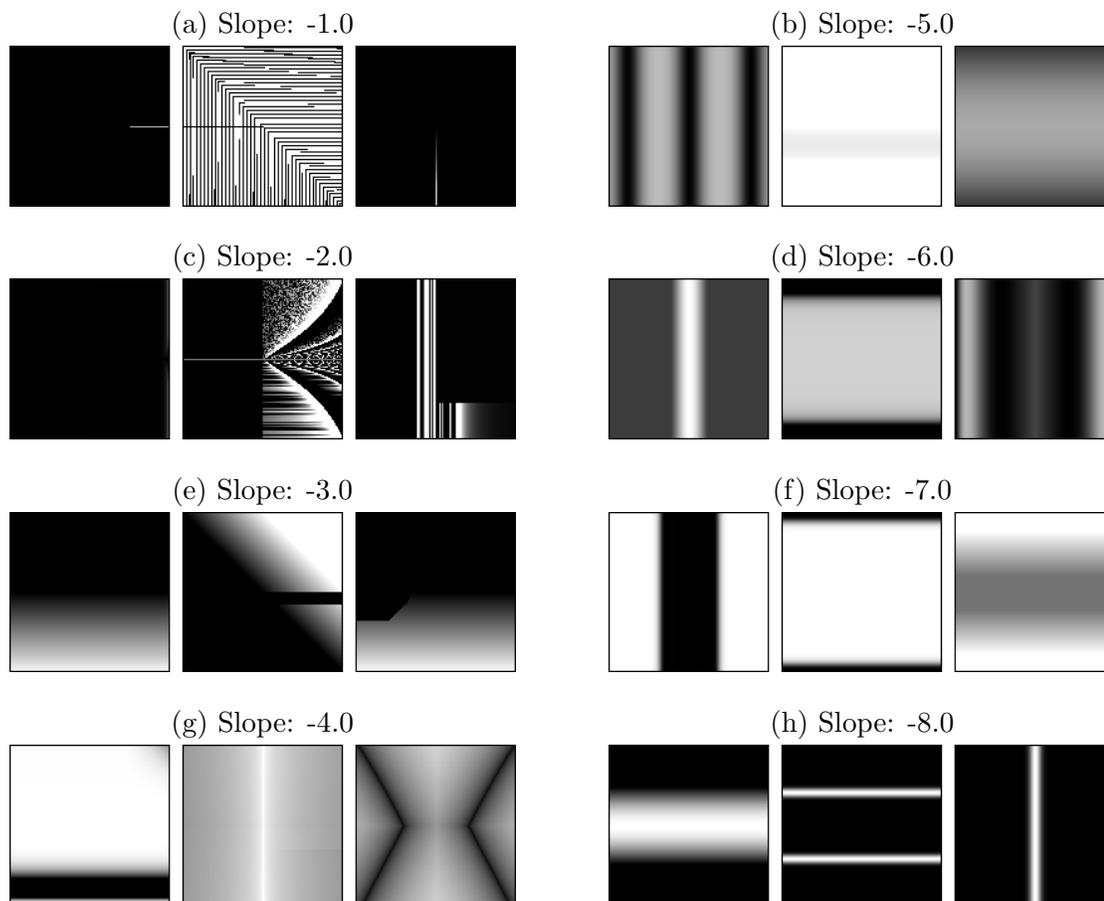
### 5.2.1 Linear Regression

With the effectiveness shown for assisting with classification and retrieval, a fitness measure which uses the slope of a linearly regressed, radially-averaged power spectra was a promising candidate for an initial exploration. Insofar that the previous literature showed an improved ability to distinguish genre by incorporating this measure, it was hoped that some spatial property capable of distinguishing these genres might emerge in our evolutionary synthesis. Being our initial foray into the exploration, it was not known what amount of fidelity could be expected in our reproduction of the target, but it was hoped that resultant images would show a clear differentiation between those evolved with different target measures.

Some substantial concerns arose early into the process of constructing the linear regression module for our GP system. While much of the earlier explored work focussed on evaluating natural images or comparatively complex art pieces, little investigation had been done into simple synthesized textures. In the process of charting the linearly averaged power spectra, and producing its regression, a transform into the log-log scale is required. While we may often expect systems capable of handling more complex individuals to also be able to handle the comparatively trivial, we run into problems when images with precisely zero power at any frequency produce

Figure 5.2: Regressed slope example images.

Target slopes were specified at regular integer intervals from  $-1.0$  to  $-8.0$ . In evaluations for a target slope of  $-8.0$ , no candidate was able to exceed a slope of  $-7.6$ , and so target slopes exceeding  $-8.0$  were found to be effectively equivalent.



infinite —or otherwise invalid— log-log scaled coefficients. In an attempt to alleviate these issues, frequencies showing precisely zero power have been removed from charting and regression, however this provides its own set of concerns. Frequencies in the spectra with zero power may incorrectly display power values averaging between their preceding and following frequencies when plotted. Removal of low-frequency coefficients may bias regression to more closely fit the higher-frequency coefficients, and vice-versa.

An even more fundamental issue with synthesizing images using this measure, beyond just retrieval, is that the single error value is too abstract. Prior to the selection of our four main, complex target images, we explored this initial measure using specified target slopes to evolve the images shown in Figure 5.2. The potential candidate images which could provide an equivalent score along this measure were too

Table 5.2: Experiment E1 Summary Table (Slope Error)

Summaries for the remainder of the chapter are produced over 30 runs. For each target, a row is included for the mean and standard deviation for each fitness objective aggregate. In the current figure, there is a single objective for slope error (distance between candidate and target) which is calculated both to find the mean slope error of a population at termination, and the best found slope error in a population at termination. The row for “mean” therein shows the mean of terminal population means, and the mean of terminal population bests across 30 runs.

|   | Target     | Agg.   | Slope Error |      |
|---|------------|--------|-------------|------|
|   |            |        | Mean        | Best |
|  | Mondrian   | Mean   | 0.13        | 0.00 |
|   |            | StdDev | 0.04        | 0.00 |
|  | Van Gogh   | Mean   | 0.09        | 0.00 |
|   |            | StdDev | 0.03        | 0.00 |
|  | Cable Ends | Mean   | 0.05        | 0.00 |
|   |            | StdDev | 0.03        | 0.00 |
|  | Flower     | Mean   | 0.20        | 0.00 |
|   |            | StdDev | 0.06        | 0.00 |

numerous. Having a fitness criteria which was too easily satisfied, and with language biases prevalent in our choice of simple mathematical operators, our GP system had a fondness of converging to simple texture results (albeit from often messy program trees).

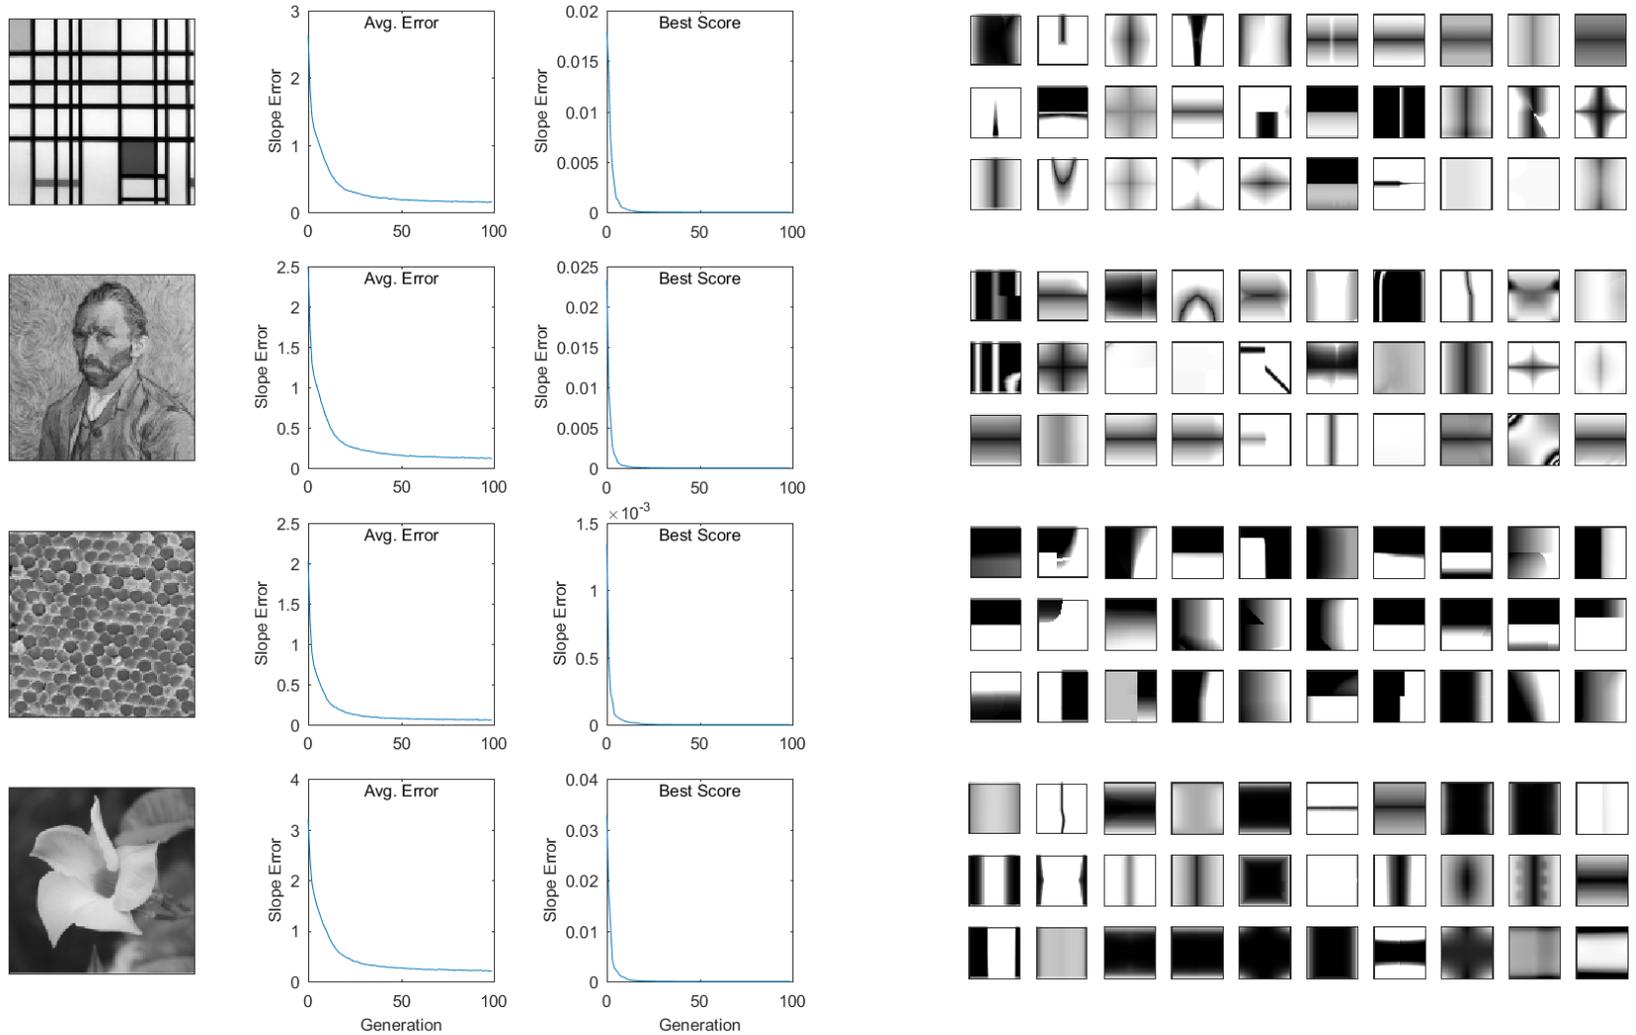
### Regressed Slope

Details pertaining to the calculation of the linear regression can be found in Section 4.1, and an overview of how we pipeline the measures’ abstractions is shown in Section 2.3.1.

Despite extolling the benefits that have been seen in applications making use of the regressed slope measure, the key information being captured is simply a rough proportion of power between higher and lower spatial frequencies in an image. This has clearly been shown to be a notable piece of information for classification, but it may be too crude for the purposes of synthesis. Where the previous research largely focuses on classification or retrieval within similar image sets, or using the regressed slope in addition to other measures, synthesis remains a challenge due to the breadth of the search space. There is an exorbitant amount of possible images that can satisfy a narrow fitness measure, many of which will likely be novel, though possibly undesirable by the user.

Figure 5.3: Experiment E1 Summary Charts &amp; Examples (Slope Error)

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 across generations (typically 100). 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.



When we nonetheless examine the effectiveness of the fitness scheme on the four main target images (shown in Figure 5.3), we see results which appear very similar to the previously fabricated target measures. For the aforementioned reasons, this was not unexpected. Rather, we see the anticipated quick convergence to a near-perfect scoring individual within the first few generations. Despite the complex target provided, our measure is too abstract; GP provides us with a high-scoring individual with extremely simple and geometric properties. From Table 5.2, we can see that the best results that were produced for each of the targets had ideal numeric scores, and were consistently produced to this quality across runs. While certain targets may slightly favour particular geometric patterns or traits, the resultant images did not appear to be generally differentiable between targets. In conjunction with the similarities found between images produced with artificially constructed slope targets, the slope measure alone is found to be insufficient in capturing any sufficient amount of spatial details.

### Regressed Slope & Offset

The next step on providing further constraint to the evolutionary process would be to also include the second coefficient in the linear regression, the line vertical offset. The possible combinations of slope and offset values prohibit us from constructing and evaluating them at even conservative interval values, and so we see the renewed need for our previous selection of experimental target images.

We had added the linear regression offset measure as a separate objective, as normalizing either measure was not practical, and a normalized sum-of-ranks fitness scheme was used for reasons outlined in Section 4.2. To clarify, we outline each objective separately in Equations 5.1 and 5.2, where slope and offset are provided from the least-squares regression.

$$Error_{slope} = |slope_{target} - slope_{candidate}| \quad (5.1)$$

$$Error_{offset} = |offset_{target} - offset_{candidate}| \quad (5.2)$$

In performing least-squares linear regression, we have seen the slope roughly describe a proportion of low and high frequencies in the image. The offset, in images with equal slope, should then estimate the comparative amount of total power in the image. However, with a log-log scaling, we see a higher density of data points lay in the higher-frequency end of our power spectra. We expect that the offset might reflect

Table 5.3: Experiment E2 Summary Table (Slope + Offset)

|   | Target     | Agg.   | Slope |      | Offset |      |
|---|------------|--------|-------|------|--------|------|
|   |            |        | Mean  | Best | Mean   | Best |
|  | Mondrian   | Mean   | 0.02  | 0.00 | 0.15   | 0.00 |
|   |            | StdDev | 0.03  | 0.00 | 0.08   | 0.00 |
|  | Van Gogh   | Mean   | 0.11  | 0.00 | 0.04   | 0.00 |
|   |            | StdDev | 0.04  | 0.00 | 0.09   | 0.00 |
|  | Cable Ends | Mean   | 0.08  | 0.00 | 0.06   | 0.00 |
|   |            | StdDev | 0.04  | 0.00 | 0.11   | 0.00 |
|  | Flower     | Mean   | 0.20  | 0.00 | 0.22   | 0.00 |
|   |            | StdDev | 0.06  | 0.00 | 0.11   | 0.00 |

an amount of power found in the higher spatial frequencies of the image. With the power spectra estimate treating positive and negative amplitude values equivalently, there could be widely varying ways for power at high frequencies to adjust, especially when less consideration is held for the lower frequencies.

With the introduction of the offset measure as a fitness objective, we can now provide enough guidance to the evolutionary system to produce more distinct sets of images. When we examine the results of the four target images in Figure 5.4, we can now begin to see some consistent traits, primarily in regards to average luminance and contrast. However, the resultant images are still extremely simple, and do not show strong spatial resemblance to their targets. Some spatial properties are being captured, but they do not appear to sufficiently align with human perception.

Much like in the previous experiment with only a slope measure, Table 5.3 shows the clear ease with which the system is able to provide numerically ideal candidates while using two regression objectives. After sufficient convergence was seen, some targets would occasionally show spikes in the amount of error for one objective while optimizing the other. The Van Gogh target showed increased noise in its offset measure during the later generations, while the flower target showed sacrifices to slope. This perhaps suggests that neither objective was inherently more difficult, but instead target dependent.

## 5.2.2 Radially Averaged Power Spectra

With the measurements in the previous two experiments being too abstract to provide sufficient detail for synthesis, our next strategy is to step out from the linear regression, and instead attempt to match the curve of the radially-averaged power spectra. While the regression could give an approximation of the proportions of high and low

Figure 5.4: Experiment E2 Summary Charts & Examples (Slope + Offset)

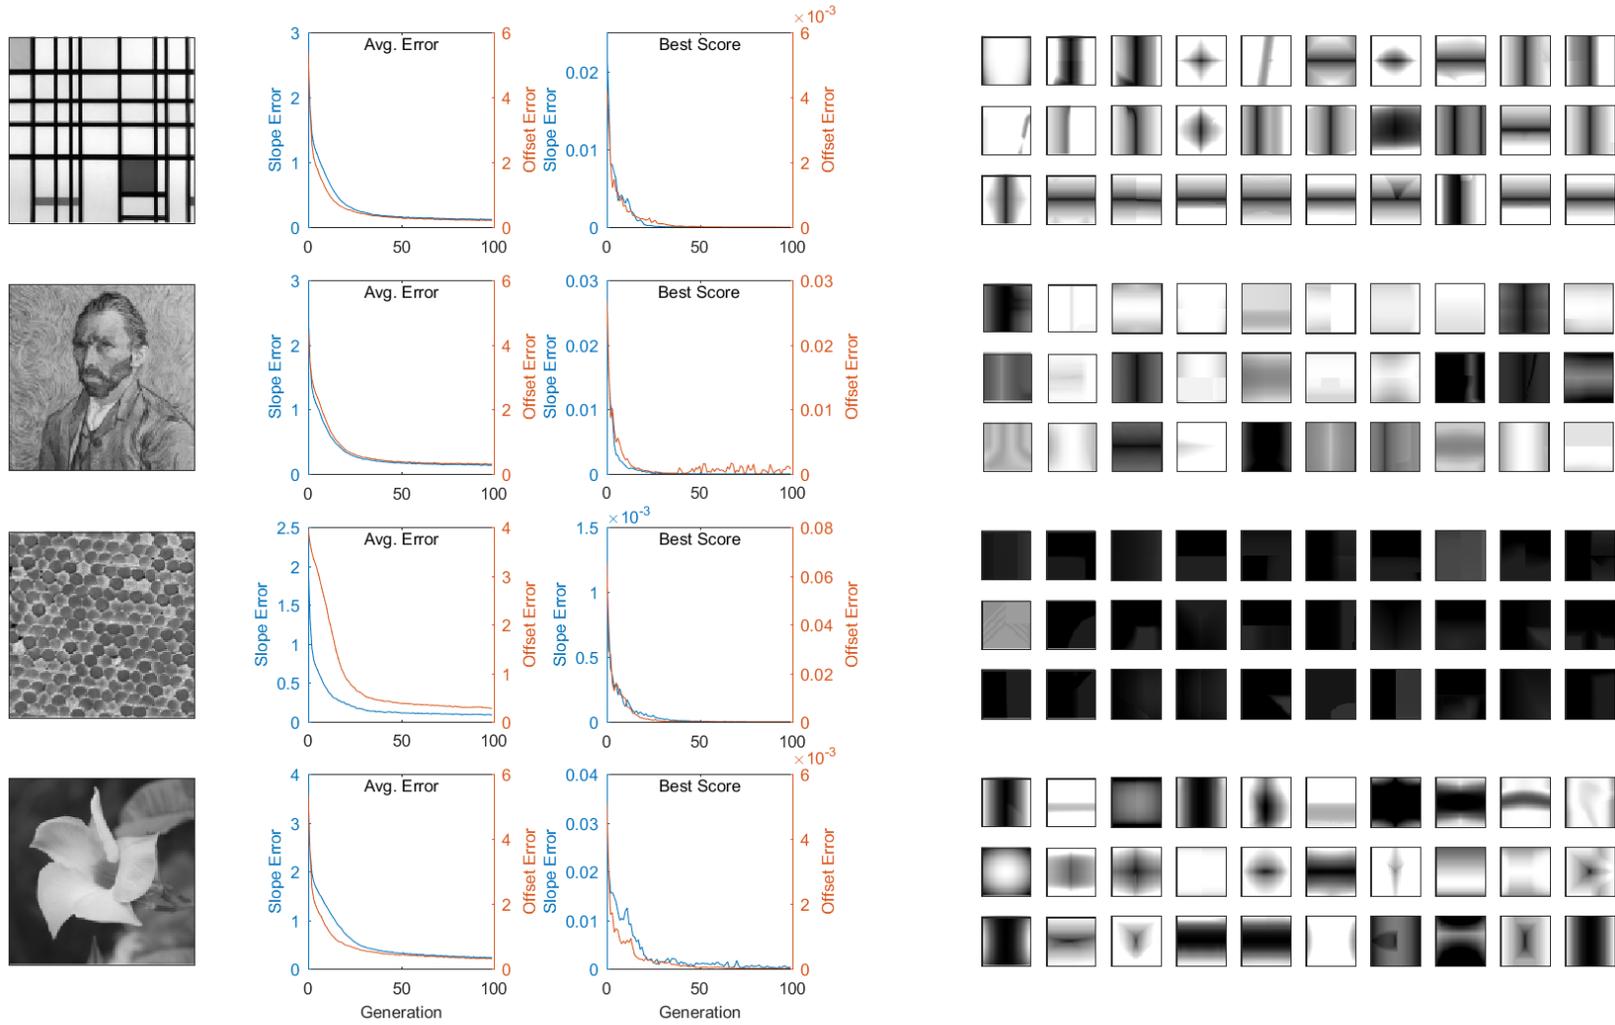


Table 5.4: Experiment E3 Summary Table (Curve Fitting)

|   | Target     | Agg.   | Curve Fit Error |         |
|---|------------|--------|-----------------|---------|
|   |            |        | Mean            | Best    |
|  | Mondrian   | Mean   | 29247078.74     | 1490.67 |
|   |            | StdDev | 16016073.04     | 516.65  |
|  | Van Gogh   | Mean   | 24723467.21     | 174.29  |
|   |            | StdDev | 15038002.79     | 242.65  |
|  | Cable Ends | Mean   | 24403723.59     | 11.53   |
|   |            | StdDev | 14825076.61     | 11.63   |
|  | Flower     | Mean   | 26113798.08     | 401.30  |
|   |            | StdDev | 23936488.33     | 479.07  |

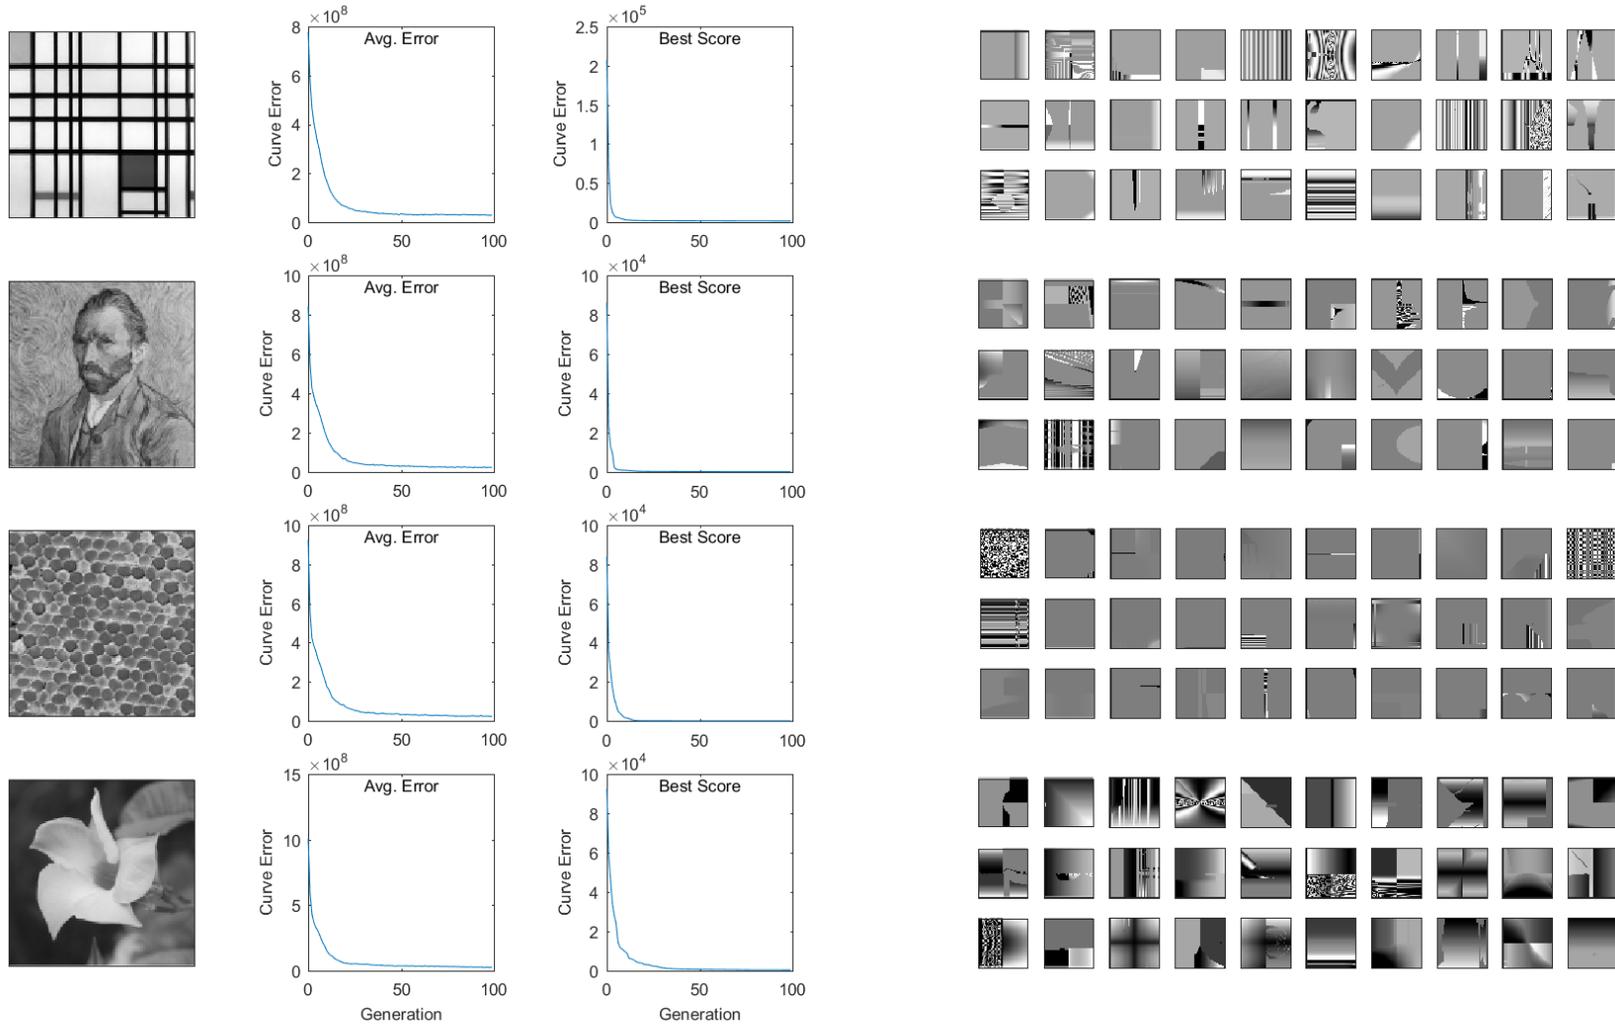
frequency components in the images, a curve fitting approach should be able to much more precisely match the power at each specific spatial frequency. As the log-log scaling (used to match the charting from existing literature) is no longer required, we can also gain the benefit of including and matching zero-powered frequencies. While a separate objective was used to handle the addition of the regression offset measure in the previous experiment, accommodating each of the 64 measured component frequencies would put us into many-objective territory, and our sum-of-ranks fitness strategy would be inadequate. Instead, we will consider a single objective from the sum of square error at each frequency. One key concern with this approach would be that most of the power in images is set about the lower frequencies (as also seen in [32]), and there is thus a high potential for bias to match the lower frequency components.

Section 4.1 provides example code for producing the radially averaged power for each frequency, from which the sum of square error can be produced using Equation 5.3. For an  $n \times n$  power coefficient matrix, we can produce the radial average of the power coefficients grouped by frequency for both the target ( $\bar{T}$ ) and candidate image ( $\bar{C}$ ), and index each frequency from 0 to  $n/2$ .

$$Error = \sum_{f=0}^{n/2} (\bar{T}_f - \bar{C}_f)^2 \quad (5.3)$$

Visually, it may not be clear if there was any improvement in comparison to the previous experiment. Figure 5.5 shows similarities between images across the various targets. While not as easily differentiable, there are some minor improvements to visual similarity with the images' respective targets. The flower target appears to have stronger, more consistent contrasts emerging from most of its evolved images.

Figure 5.5: Experiment E3 Summary Charts & Examples (Curve Fitting)



Beyond that, the Mondrian target has a few images with themes of spaced lines, the Van Gogh has a few images with some vague aspect of a profile, and a pair of images evolved from the Cable Ends target has a fine granular noise-like texture. However, these appear to be exceptions, not the typical result.

It would seem that we are capturing more useful information, but it is not clear if what we capture is sufficient, or if our measure is too difficult for the evolutionary system to use as guidance. From Table 5.4, we see that the average population error is orders of magnitude larger than the error of the best individual, but this may not be unexpected due to reproduction operators having the potential to cause huge changes to the lower —more powerful— frequencies. Rather, the charting in Figure 5.5 displays a suitable convergence.

Returning to Table 5.4, we can see that the error levels of the best evolved images are orders of magnitudes less than the final population mean. Comparing the difference in error between the best found individuals for each target, it would seem that certain targets are inherently easier for our system to satisfy. Worrying, however, is the low error found with the Cable Ends target. While two images produced for that target had some remote similarity, most appeared quite visually dissimilar. This strongly suggests that some visual component is not being captured by our measure. It is quite possible that the radial averaging is providing leniency; harmonic frequencies could be missed if recreated at different radial angles (and thus in the same bin for radial averaging). We will also later adjust our fitness metrics and incorporate a measure for component frequency phase, but not before refining the capabilities from a purely power-based approach.

### 5.2.3 Power Coefficients

The next step in our set of experiments is to match power coefficients. While a perfect match of coefficients would diminish the purpose of producing new, novel images, the absence of capturing any phase information still leaves us with a plentiful amount of possible evolutions. Concerns that the previous curve fitting approach could be obscuring certain frequency interactions further warrants a less abstracted fitness measure. We will simply employ a single objective, measuring the sum of squared error across each of the  $128^2$  power coefficients. (To improve computational efficiency, half of the coefficients may be ignored due to the symmetrical nature of FFT on non-imaginary input, and the resultant power spectra.) This measure is described in Equation 5.4, assuming an  $n \times n$  power coefficient set, with  $T$  belonging

Table 5.5: Experiment E4 Summary Table (Coefficient Error)

|   | Target     | Agg.   | Coefficient Error |           |
|---|------------|--------|-------------------|-----------|
|   |            |        | Mean              | Best      |
|  | Mondrian   | Mean   | 40801723.31       | 687983.82 |
|   |            | StdDev | 19789631.85       | 84851.29  |
|  | Van Gogh   | Mean   | 29554765.71       | 5167.94   |
|   |            | StdDev | 18628699.02       | 2016.67   |
|  | Cable Ends | Mean   | 25481751.48       | 1855.43   |
|   |            | StdDev | 15398753.35       | 85.77     |
|  | Flower     | Mean   | 24287423.34       | 28761.09  |
|   |            | StdDev | 13255031.22       | 16807.58  |

to the target and  $C$  the compared candidate. Much like the curve fitting measure, we expect to see a bias toward lower frequency coefficients.

$$Error = \sum_{y=1}^n \sum_{x=1}^n (T_{x,y} - C_{x,y})^2 \quad (5.4)$$

Visually, from Figure 5.6, we see mixed results across the targets. The Mondrian targets appears to show improved consistency with recreating lines, and the flower target continues to show strong light and dark contrasts. The other two targets appear to have worsened visual performance, despite showing the strongest results numerically, as in Table 5.5. Despite our preconceived expectation that the Mondrian target should be the easiest among initial targets to recreate with our system, and despite it being a target with many visually similar evolved candidates, it is the worst scoring of the targets for this experiment. Additionally, the visual disparity between the Cable Ends target and its evolved candidates despite its numeric performance is curious.

When we examine some of the evolved candidates and compare them to the targets (Figure 5.7), our observations are completely opposite to the numeric scores above. Despite the low amount of error produced with the Cable Ends target, the power coefficient display is substantially different between the target and its visually dissimilar evolved candidates. Further, some of the more visually faithful candidates from the Mondrian target have similar power coefficient displays despite their numeric error scores. This is actually a welcome sanity check. These observations would suggest that the first few coefficients provide far too heavy a bias to produce meaningful results, and that some other approach will be needed to overcome this bias.

Figure 5.6: Experiment E4 Summary Charts & Examples (Coefficient Error)

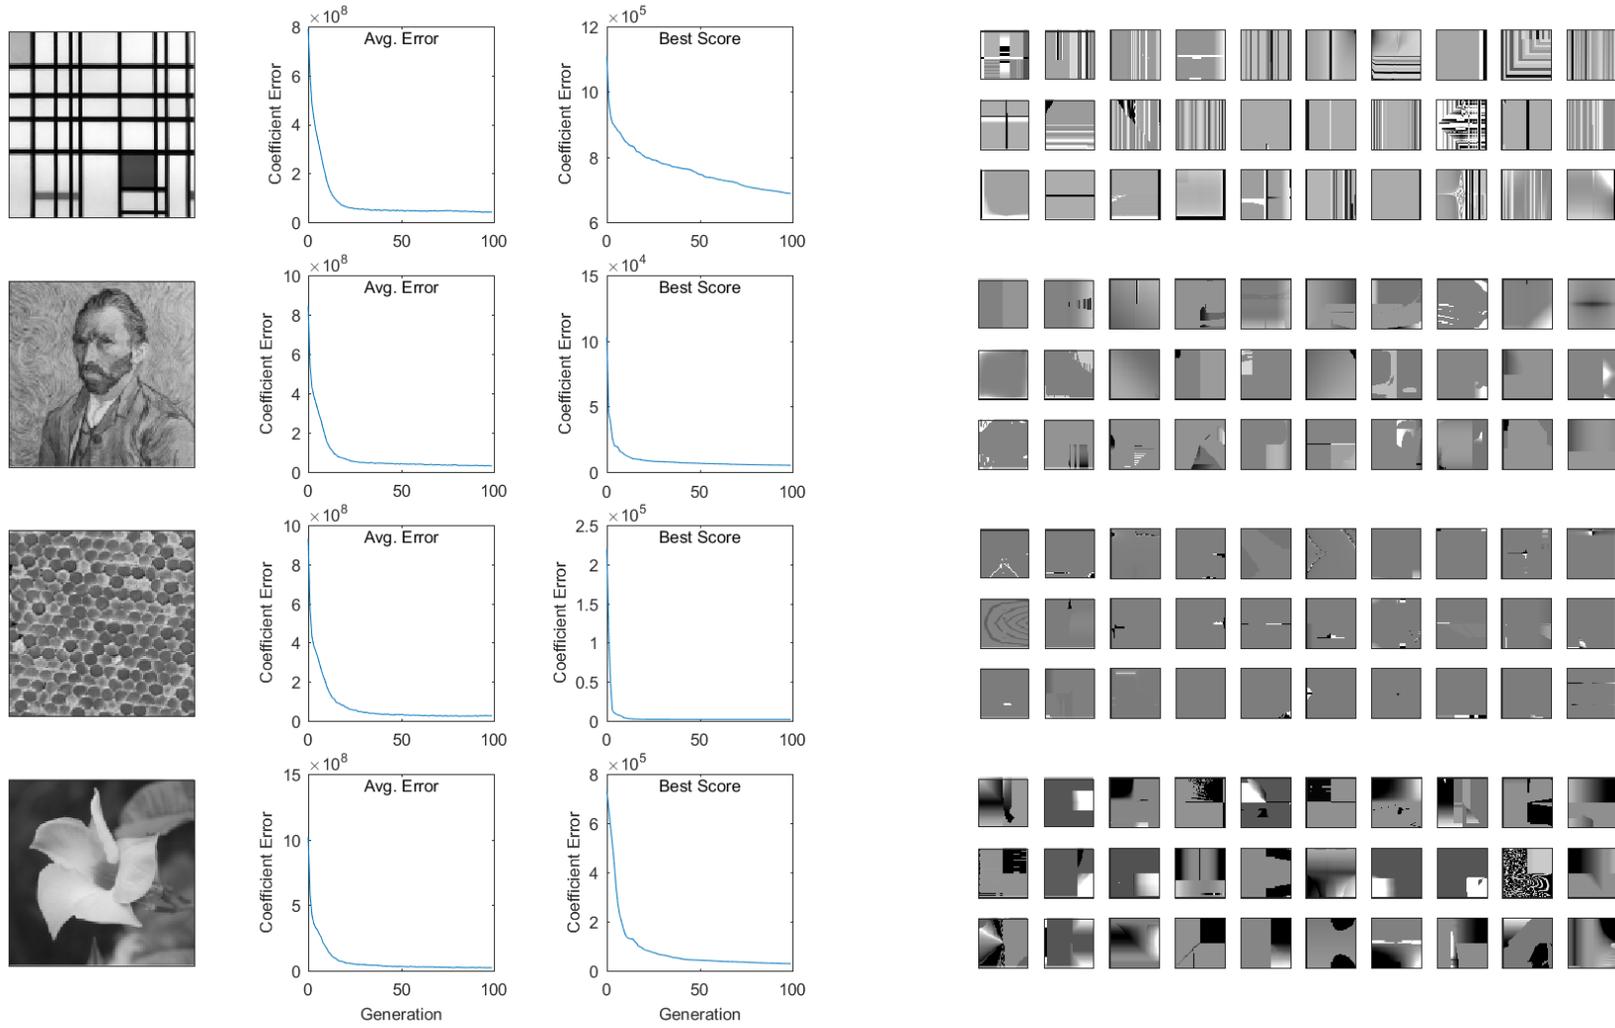
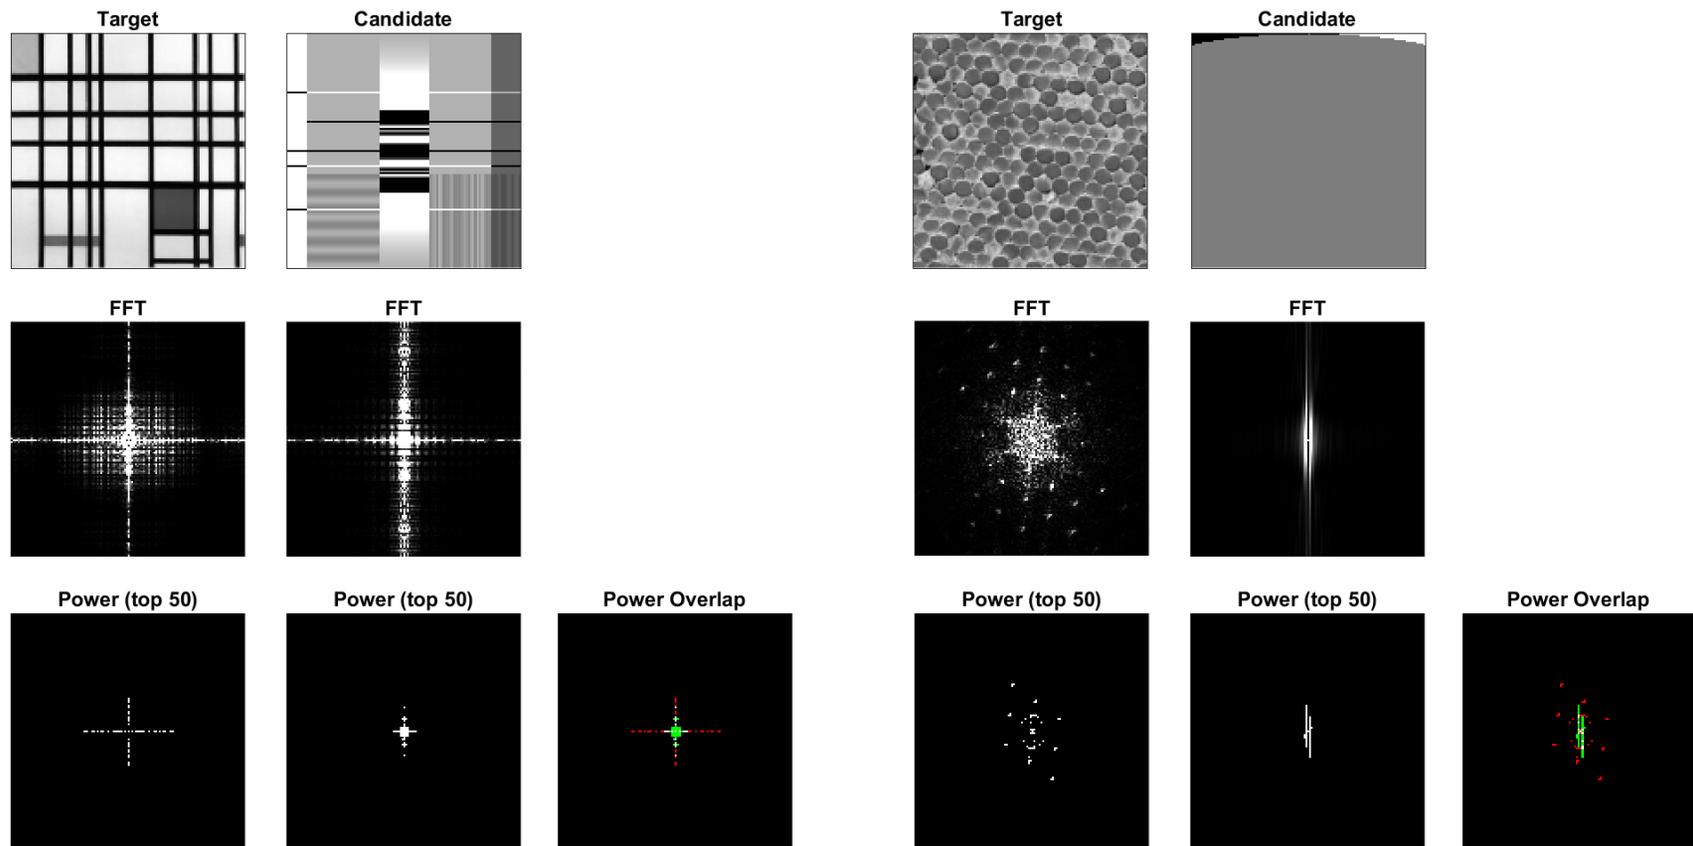


Figure 5.7: Experiment E4 Target &amp; Candidate Comparison

Evolved images are compared alongside their target image. The second row of each set shows the complete set of power coefficients, where the third row filters displayed coefficients to the top 50. Overlap between the filtered top 50 is shown in the style of a confusion matrix; common positions are white, positions missed from the target are red, and positions in the candidate but not target are green. This shows the relative ranked status of the top 50 coefficients.



### 5.3 Filtering Relevant Coefficients

From our attempts at guiding texture evolution through power coefficient error, we have gained some insights on how few coefficients can heavily influence the fitness measure. It follows that we will need to explore various ways of more fairly scoring similarity between power coefficient sets. We will craft and evaluate new fitness measures which play with ideas of extracting smaller sets of coefficients, use rank distance in lieu of error distance, separate more powerful coefficients to their own objectives, and begin to incorporate aspects of phase. An overview of the experiments considered in this section can be found in Table 5.6 below.

Table 5.6: Overview of Section 5.3 Experimental Variations

| Label | Measure | K   | Language | Tables | Figures | Notes   |
|-------|---------|-----|----------|--------|---------|---|
| J1    | J1      | 50  | Base     | 5.7a   | 5.9     | Top K truncated, quantized error (position mismatch)                    |
|       | J1      | 150 | Base     | 5.7b   | 5.9     |   |
| J2    | J2      | 50  | Base     | 5.8a   | 5.10    | Rank distance error   |
|       | J2      | All | Base     | 5.8b   | 5.10    |   |
| J3    | J3      | 50  | Base     | 5.9a   | 5.13    | J1 with integrated phase error (linear)                                 |
|       | J3      | 50  | Base     | 5.9b   | 5.13    | J1 with integrated phase error (squared)                                |
| K1    | K1      | 50  | Base     | 5.10   | 5.14    | Top K, non-quantized power and phase error                              |
| K2    | K2      | 50  | Base     | 5.11   | 5.15    | K1, with separate objectives for power and phase                        |
| K3    | K3      | 50  | Base     | 5.12a  | 5.16    | K2, with the 0 <sup>th</sup> coefficient separated to a third objective |
|       | K3      | 150 | Base     | 5.12b  | 5.16    |   |

#### 5.3.1 Coefficient Extraction from Jacobs *et al.*

An initial survey into extracting key coefficients from a 2D DFT decomposition did not yield many useful results. Most of the found articles pertain to compression of sparse transforms, and while doubtless useful for other applications, did not provide much in means of determining spatial saliency. The common feature between papers [32][36], and our existing understanding with image composition, was in the prioritization of

higher power coefficients. However, this does not assist us with our previously found issue of coefficient bias.

A strategy for frequency analysis coefficient isolation was found among works in image retrieval using wavelets. The paper on *Fast Multiresolution Image Querying* by Jacobs *et al.* [45] outlines a technique for efficient retrieval of images by comparing how well a set of extracted wavelet coefficients and their positions match. The technique had been adapted for use with texture synthesis previously [43][44] with some level of reported success. In the initial paper by Jacobs *et al.*, the calculation of the querying metric from the transformed coefficients is rather straightforward and efficient. First, the set of coefficients is “truncated” by zeroing all but the top  $K$  greatest absolute value coefficients. Following this is a “quantization”, setting all non-zero components into their sign of  $\{-1, +1\}$ . The total error between images is then the sum of differences between each truncated, quantized coefficient position. The paper also plays with the idea of binning coefficients by position and using a weighted sum instead for its final calculation.

A promising observation from the paper regards the beneficial effect of the quantization scheme. It was seen that this quantization improved the ability of metric, and despite losing much of the precise data from the coefficients, “the mere presence or absence of such features appears to have more discriminatory power for image querying than the features’ precise magnitudes” [45]. This is wonderfully in line with our own previous observations regarding a small number of high powered coefficients overwhelming the ability of our metric to fairly guide evolution.

A few further considerations will need to be made before attempting a similar scheme using Fourier transforms instead of wavelets. A quantization to  $\{-1, 0, +1\}$  is not as meaningful in the context of a Fourier transform, where power coefficients are strictly positive. Amplitude coefficients may hold negative values, but also show negative symmetries, and can change sign when set with opposite phase. We can truncate coefficients as per the paper, but the solution we have attempted will instead quantize all remaining values simply to 1 (an example of this can be seen in Figure 5.8). This effectively turns the score into a count of how many positions share a top  $K$  coefficient. As any position giving an error when observing the top  $K$  locations from the target necessarily means there is an additional error found when observing the top locations in the candidate, we can avoid double counting by only observing the errors in the targets’ top positions.

While a wavelet decomposition requires further choices for wavelet type, decomposition type, and basis normalization schemes, Fourier compositions are constrained

Figure 5.8: Truncation and Quantization Example

A simplified example with superscript noting the relative rank of the coefficient position. Positions with error have been displayed in red.

$$\begin{bmatrix} 0 & +2 & +4 & 0 & 0 \\ +1 & +1 & +5 & -1 & +1 \\ -1 & -5 & +20 & +5 & +1 \\ -1 & +1 & -5 & -1 & -1 \\ 0 & 0 & -4 & -2 & 0 \end{bmatrix} \xrightarrow{K=9} \begin{bmatrix} 0 & 1^8 & 1^6 & 0 & 0 \\ 0 & 0 & 1^2 & 0 & 0 \\ 0 & 1^4 & 1^1 & 1^5 & 0 \\ 0 & 0 & 1^3 & 0 & 0 \\ 0 & 0 & 1^7 & 1^9 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & +1 & +9 & 0 & 0 \\ +1 & +1 & +8 & +3 & 0 \\ -1 & +2 & +19 & -2 & +1 \\ 0 & -3 & -8 & -1 & -1 \\ 0 & 0 & -9 & -1 & 0 \end{bmatrix} \xrightarrow{K=9} \begin{bmatrix} 0 & \mathbf{0} & 1^2 & 0 & 0 \\ 0 & 0 & 1^4 & \mathbf{1^6} & 0 \\ 0 & 1^8 & 1^1 & 1^9 & 0 \\ 0 & \mathbf{1^7} & 1^5 & 0 & 0 \\ 0 & 0 & 1^3 & \mathbf{0} & 0 \end{bmatrix}$$

but simplified. We do, however, still need to choose a  $K$  value to determine the size of our coefficient truncation. Jacobs *et al.* had found values of 40 to 60 to perform well with their image retrieval data sets, with a slightly increased  $K$  value showing leniency (improved matches) when retrieving rougher, hand-drawn recreations of the queried image. In our selection of a suitable  $K$  value, we considered possible reconstructions of the target images which had power removed from all but the top  $K$  positions. A wide range of values for our selection of  $K$  were found to produce similar recreations of the target when maintaining similar phase values, but prominent recreations seemed to begin to form in the range of  $K = [50, 150]$ . Target image reconstructions at various  $K$  choices can be seen within Figure A.1 in Appendix A.

### Initial Adaptation

With a proposed countermeasure to the few domineering coefficients, our next experiment will be attempted using our adaptation of the Jacobs *et al.* metric. Continuing without changes from our base system, and maintaining the same basic GP language, we will explore evolutions with the selections of both  $K = 50$  and  $K = 150$ .

Initial review of the evolved solutions in Figure 5.9 looks rather rough for both choices of  $K$ . With  $K = 50$ , the Flower target has some minor aspects appear in its solutions, and the Cable Ends target has a somewhat consistent display of partial stripes —seemingly proportioned to the width of a cable strand end in the target. When we follow up by referencing Table 5.7, we see that best results for the flower target at  $K = 50$  had an average of  $7.3/50 \approx 15\%$  mismatched coefficients in its top ranks. These images produced are a fair bit rougher, though not entirely

Table 5.7: Experiment J1 Summary Table

(a)  $K = 50$ 

|  | Target     | Agg.   | Coef. Mismatch |       |
|--|------------|--------|----------------|-------|
|  |            |        | Mean           | Best  |
|   | Mondrian   | Mean   | 32.46          | 15.17 |
|  |            | StdDev | 1.03           | 1.76  |
|   | Van Gogh   | Mean   | 30.41          | 11.80 |
|  |            | StdDev | 0.83           | 1.16  |
|   | Cable Ends | Mean   | 40.79          | 30.27 |
|  |            | StdDev | 1.06           | 2.66  |
|  | Flower     | Mean   | 27.81          | 7.30  |
|  |            | StdDev | 0.89           | 1.15  |

(b)  $K = 150$ 

|   | Target     | Agg.   | Coef. Mismatch |       |
|---|------------|--------|----------------|-------|
|   |            |        | Mean           | Best  |
|  | Mondrian   | Mean   | 53.28          | 18.37 |
|   |            | StdDev | 2.16           | 1.87  |
|  | Van Gogh   | Mean   | 69.52          | 34.33 |
|   |            | StdDev | 3.02           | 2.12  |
|  | Cable Ends | Mean   | 116.26         | 88.97 |
|   |            | StdDev | 2.41           | 3.50  |
|  | Flower     | Mean   | 63.90          | 27.37 |
|   |            | StdDev | 2.62           | 1.97  |

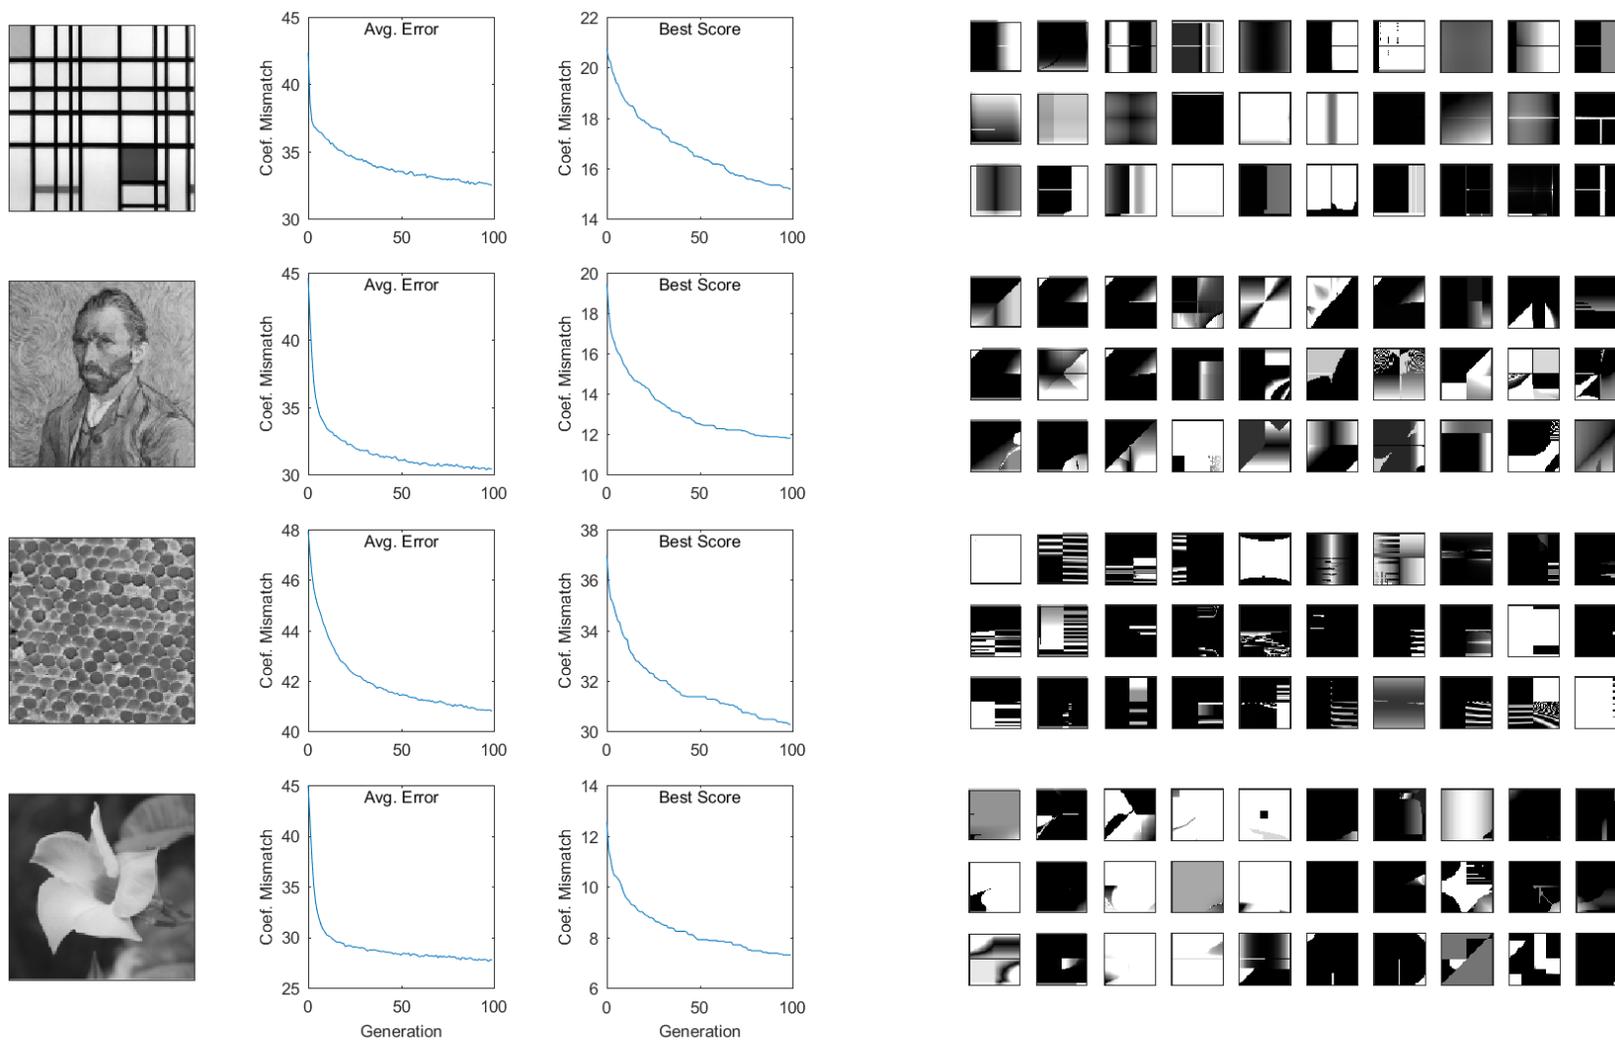
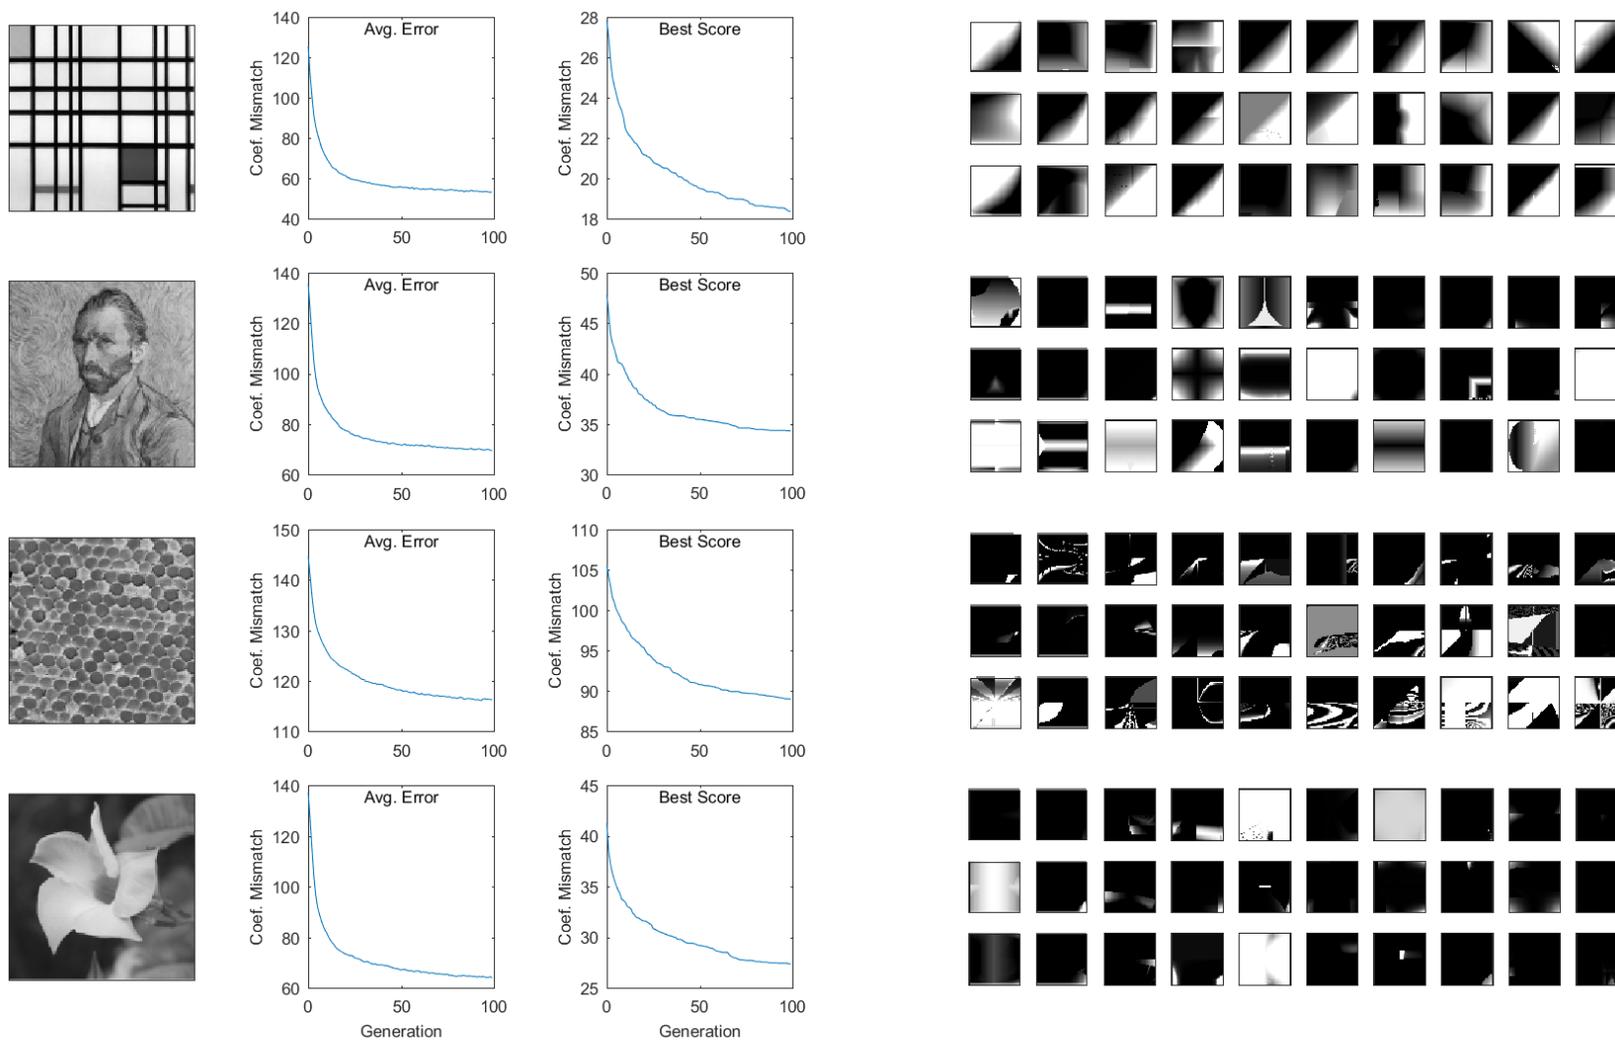
Figure 5.9: Experiment J1 Summary Charts & Examples;  $K = 50$ 

Figure 5.9: Experiment J1 Summary Charts & Examples;  $K = 150$ 

unimaginable from the types of images we have seen at similar  $K$  values during image reconstruction (Appendix A, Figure A.1). While still very far from the target image, the candidates produced from the Cable Ends image do have some minutely familiar spatial properties, and accomplish this despite having over 60% mismatch on average. The remaining targets have seen some increased contrast in comparison to previous experiments, but altogether have not shown substantial visual improvements.

When we consider the change to  $K = 150$ , the amount of error has strictly increased, though not uniformly nor linearly. Despite tripling the count of coefficients remaining after truncation, some error scores have approximately tripled, but others have increased further, and some —notably the Mondrian target — have barely risen at all. Further, despite the Mondrian target showing a capability of matching higher coefficients, we have seen a substantial visual drop at  $K = 150$ . The increased  $K$  value appears to give a consistent visual degradation across all of the target images. With the Mondrian target, it appears that almost all found solutions have adopted a diagonal contrast as it is presumably much easier to incorporate some of the lower power coefficients near the diagonal centre. Choice for  $K$  will need to be balanced between providing the necessary level of detail for the target, without providing too many options for the result to be misguided.

We will shortly see further trade-off for choice of  $K$  values when considering the integration of phase error in our metrics.

### Rank Distance Adaptation

In considering other methods to reduce the overwhelming influence of the few largest coefficients, inspiration was taken from other parts of our GP system to consider difference in rank. With the suggestion from Jacobs *et al.* that placement in the top positions are more effective discriminators than the actual coefficient values, we can further consider that relative powers amongst the component frequencies is an alternative worth consideration. We go forth with the idea that frequencies ranked higher amongst others in the target should also rank higher amongst others in the candidate image. By summing the total difference in rank number for each coefficient between target and candidate, we have a potential measure for disparity in how relatively important each component frequency is.

Conceivably, this new measure may not require the truncation procedure used by Jacobs *et al.*, however, we will consider the cases of both  $K = 50$  and  $K = 128^2$  (all coefficients). Checking the rank error on all positions may provide more total precision in the ideal case, but could also easily result in poor-looking results due

Table 5.8: Experiment J2 Summary Table

(a)  $K = 50$ 

|   | Target     | Agg.   | Rank Error |          |
|---|------------|--------|------------|----------|
|   |            |        | Mean       | Best     |
|  | Mondrian   | Mean   | 9949.74    | 942.40   |
|   |            | StdDev | 4901.56    | 89.51    |
|  | Van Gogh   | Mean   | 8299.26    | 718.53   |
|   |            | StdDev | 3288.52    | 91.19    |
|  | Cable Ends | Mean   | 28716.72   | 11693.40 |
|   |            | StdDev | 7430.03    | 2575.71  |
|  | Flower     | Mean   | 8919.63    | 446.20   |
|   |            | StdDev | 6692.69    | 116.23   |

(b) All Coefficients

|   | Target     | Agg.   | Rank Error  |             |
|---|------------|--------|-------------|-------------|
|   |            |        | Mean        | Best        |
|    | Mondrian   | Mean   | 36856289.46 | 32323384.53 |
|   |            | StdDev | 2058630.46  | 742917.26   |
|   | Van Gogh   | Mean   | 58185631.10 | 53110794.60 |
|   |            | StdDev | 1817498.60  | 218671.21   |
|  | Cable Ends | Mean   | 54830748.34 | 49263217.73 |
|   |            | StdDev | 1805632.30  | 130566.58   |
|  | Flower     | Mean   | 44611553.74 | 37384364.67 |
|   |            | StdDev | 2117072.33  | 437745.62   |

to the many very low powered (effectively non-visible) coefficients happening to be correctly proportioned; an effective bias for the low-powered frequencies.

From Figure 5.10, the idea of using the rank of all  $128^2$  coefficients can be discarded. Candidate populations appear to have converged with very high remaining rank error. With the extremely simplistic resultant textures, we can see that the fitness measure was not able to provide strong enough guidance, where the population prematurely converged after tweaking the ranks of a few mid rank coefficients (still very low power and largely undetectable) and falling into a local optimum.

When we limit the ranking to the top  $K = 50$  coefficients, we compare ranks of the top 50 positions in the target image to the ranks at that respective position in the candidate. From Table 5.8, most targets performed respectably, with the noted exception of the Cable Ends target. Examining the location of the larger power coefficients in the target (Appendix A, Figure A.1), it would seem that this may be a difficult target for the measure due to the uncommon location of some of its higher-powered harmonics far off on the diagonals. The results with other targets seem to

Figure 5.10: Experiment J2 Summary Charts & Examples;  $K = 50$

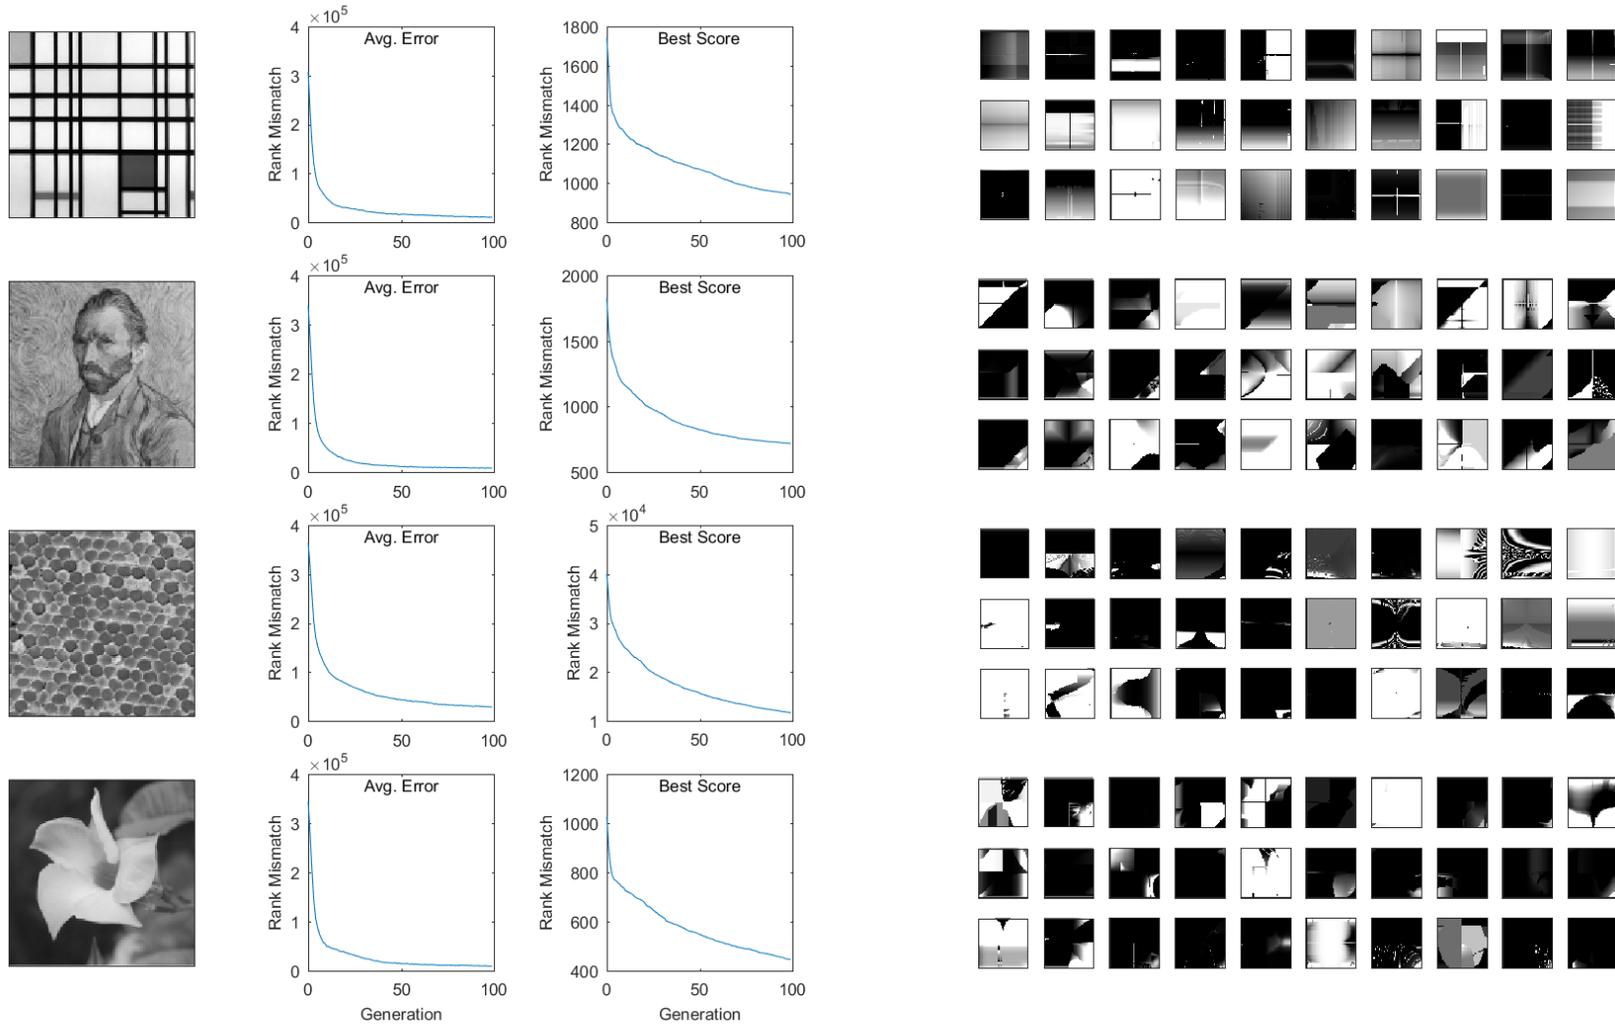
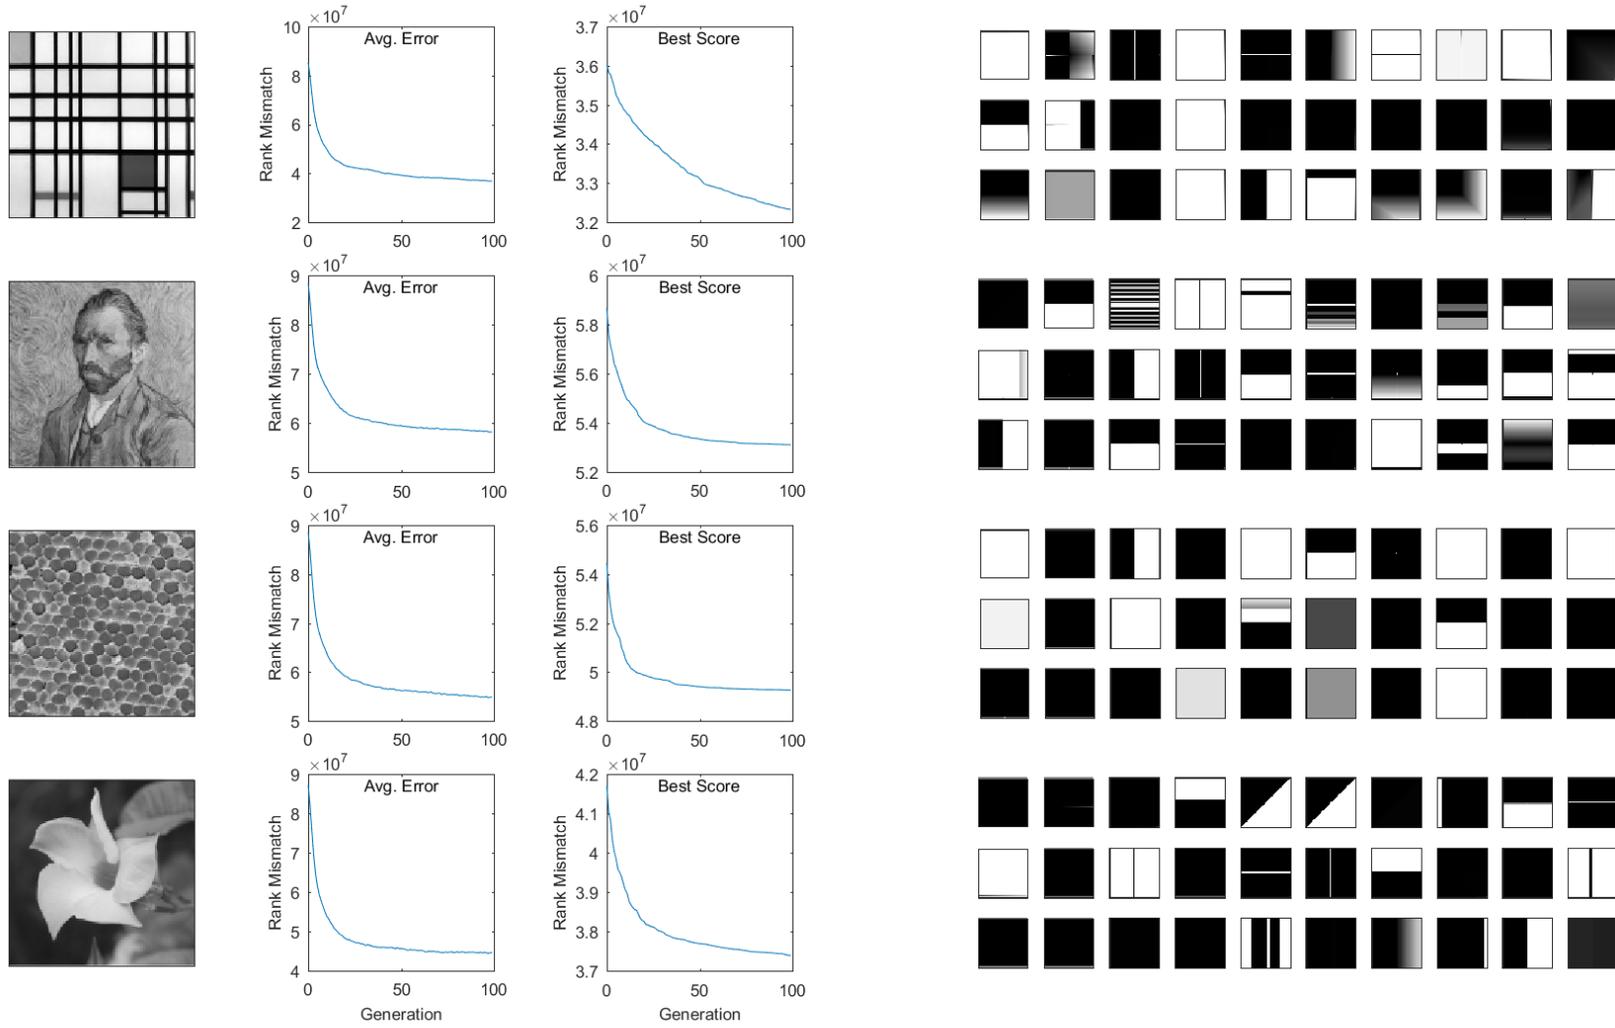


Figure 5.10: Experiment J2 Summary Charts & Examples; All Coefficients



be not too dissimilar from those produced by the adapted Jacobs *et al.* approach.

### 5.3.2 The Importance of Phase

A critical difference between the original metric put forth from Jacobs *et al.* using wavelets, and our adaptation with Fourier transforms, is the inherent removal of any spatial localization in our frequency analysis. Fourier analysis has a higher precision in measuring component frequencies throughout an entire image, yet wavelets are able to localize where power was seen in the source image based on the coefficient indices. Thus, when the measure by Jacobs *et al.* was truncating coefficients, they were capturing both some degree of frequency and also positions in the image where the frequencies manifest.

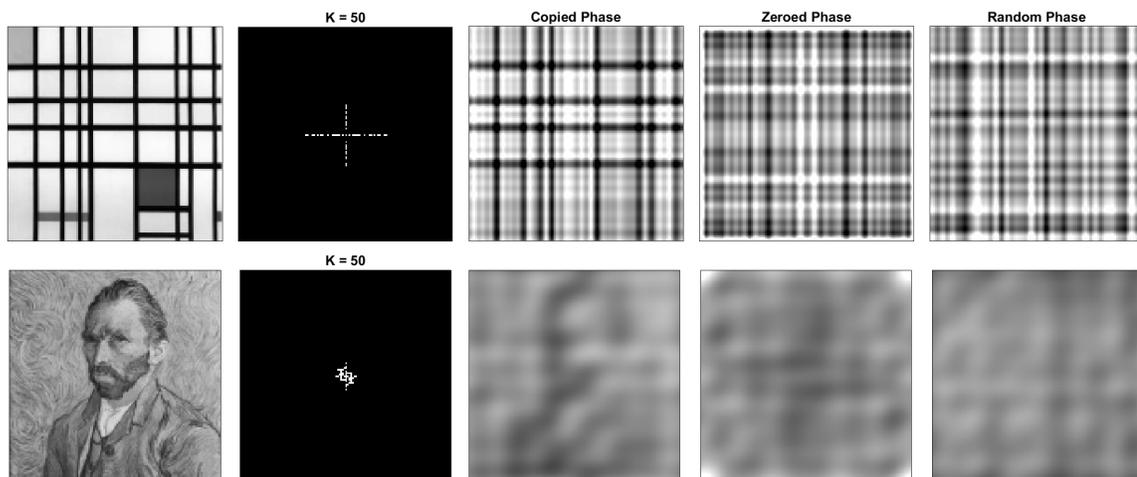
Up until now, our proposed fitness measures have attempted to produce images with similar amounts of component frequencies. When measuring coefficients, their index and position (the radial angle of the coefficient from centre) has encouraged evolution of component frequencies with similar placement, but we have largely overlooked how these component frequencies should be offset and overlap. The other key aspect of a Fourier transform, the phase component, has not been taken into account for any measure. By reincorporating this measure into our fitness schemes, we may provide some further constraints on the location of where the component frequencies crest.

#### Examining Potential Gains

With the suspicion that our lack of capturing any aspect of phase has begun to limit the extent of what we can guide through evolution, we re-examine our targets while considering what they may look like if their component frequencies had alternative phases. In Figure 5.11, we have highlighted some reconstructions of target images while limiting coefficients to their top  $K = 50$ , and providing alternative phase values. Reconstructions with various phase selections and coefficient truncations can be seen for all targets in Appendix A, Figure A.1. While there is sufficient power to begin recreating the images in the most ideal circumstances (phase retained from the original target), when we consider the case of alternative, random selection for phase angles, images can become anything from heavily varied to completely unrecognisable. These images—which may indeed show no human-recognizable similarities—would be considered ideal solutions by our existing fitness schemes. Clearly there is a need to incorporate phase into our fitness measures for more reliable human recognition. Some

Figure 5.11: Reconstructing Target Images With Varied Phase

A pair of targets are 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 positions, we can see the most salient positions alongside the target in the second column. The third column is the inversed FFT reconstruction with power limited to these truncated positions. We also show reconstructions 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.



targets, like the Mondrian image, still carry many similar traits when ignoring phase, and would instead require increased precision to which coefficient positions contain high power measures. This suggests that targets may have varying requirements for precision in phase.

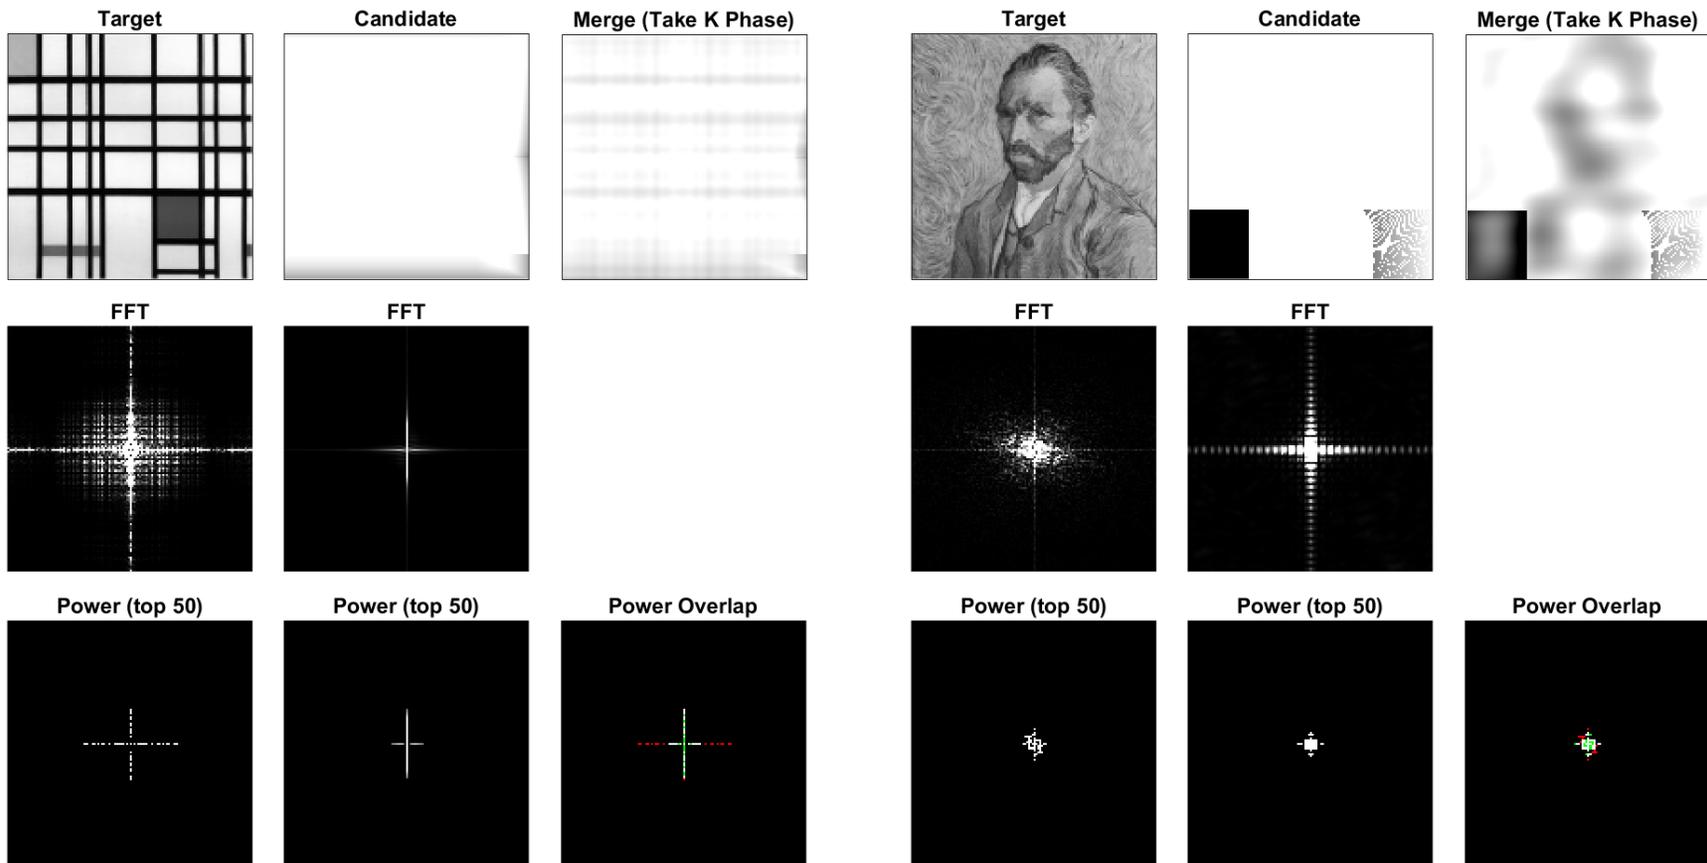
After observing the wide variety of images which could potentially satisfy our fitness measure, we also re-examine a number of previously evolved candidate images which were notable only for their poor visual resemblance to the target. We see that these candidates were often not scored among the outliers, but actually held reasonably similar top power coefficient positions. Figure 5.12 reconstructs a pair of images which we identified as having low visual similarity to their targets. By reconstructing the images using their targets' phase angles, we're given an optimistic look at the type of images that could be expected if phase was were to be successfully guided in evolution.

### Including Phase In Fitness

With the discovery that our previous measures can produce images with reasonable approximations to the targeted frequencies, we will now begin to also score phase

Figure 5.12: Reconstructing J1 Results With Target Phase

Target images and an evolved candidate are compared to observe truncated power coefficient mismatch and an optimistic estimation of results if phase angles were replicated. The previous Figure 5.7 describes the displayed power coefficients. We now also show an image reconstruction in the first row, last column, which combines the candidates power coefficients with the phase angles of the target image. The compared (poor) candidates show noted improvement when phase is adjusted.



information for further evolutionary guidance. We return to our adapted Jacobs *et al.* approach —or, top  $K$  mismatch— and will consider the difference of phase angle for those top  $K$  positions. Being mindful that phase error should wrap about  $2\pi$ , the maximum difference in phase angle should be  $\pi$ . We normalize the phase error to  $[0, 1]$  and square it for each of the top positions, using this error to slightly penalize the top matching positions if they are out of phase. So as to not completely remove the benefit of having a top ranked frequency, even if completely out of phase, we will limit the phase error by scaling it to  $[0, 0.5]$ . Effectively, for each of the target’s top  $K$  ranked coefficient positions, a penalty of 1.0 is added if the position is not also ranked among the candidate’s top  $K$ , and if it is, there could be a penalty of up to 0.5 based on phase mismatch.

$$Error = \sum_{i=1}^K \begin{cases} 0.0 + \frac{1}{2} [\pi^{-1} \Delta(\theta(T, T_i), \theta(C, T_i))]^2 & , T_i \in C \\ 1.0 & , T_i \notin C \end{cases} \quad (5.5)$$

$$\Delta(a_1, a_2) \equiv (a_1 - a_2) \pmod{\pi} \quad (5.6)$$

We encode this in Equation 5.5: where  $T$  and  $C$  are the truncated set of coefficient positions for the target and candidate as ordered by power. We will have  $\Phi(V, p)$  and  $\theta(V, p)$  return the power and phase angle respectively of the coefficients (vectors) in set  $V$  corresponding to coordinate  $p$ . With a slight abuse in notation, we will denote the coordinates of the  $i^{th}$  ranked position (by power) of a coefficient set as  $S_i$ .

After further adapting the Jacobs *et al.* measure to include the additional phase mismatch penalty, results are slightly underwhelming, but also show some slight improvements and promise. The candidates generated from these runs shown in Figure 5.13 are still quite rough, and appear largely similar to those produced by experiment J1 with the phase-less coefficient mismatch measure. However, there are a few subtle but notable differences. With the Van Gogh target, we see a more prominent triangular profile shape from the bottom to the centre of the image (optimistically, resembling the torso and edge of the shoulders). The candidates from the flower target seem to be recreating a brighter object emerging from the darker edges with some consistency. It was initially anticipated that the Mondrian target would immediately benefit from the inclusion of phase information, and so it is unfortunate that no particular differences could be recognized from its candidates produced without phase information. We could remain slightly hopeful as, for the Mondrian target specifically, the population has yet to fully converge despite showing similar numeric

Table 5.9: Experiment J3 Summary Table

(a)  $K = 50$ , Linear Error

|  | Target     | Agg.   | Combined Error |       |
|--|------------|--------|----------------|-------|
|  |            |        | Mean           | Best  |
|   | Mondrian   | Mean   | 24.39          | 20.43 |
|  |            | StdDev | 1.51           | 2.03  |
|   | Van Gogh   | Mean   | 18.91          | 15.40 |
|  |            | StdDev | 1.39           | 1.19  |
|   | Cable Ends | Mean   | 34.04          | 30.87 |
|  |            | StdDev | 1.99           | 2.46  |
|  | Flower     | Mean   | 17.60          | 13.95 |
|  |            | StdDev | 1.38           | 1.26  |

(b)  $K = 50$ , Squared Error

|   | Target     | Agg.   | Combined Error |       |
|---|------------|--------|----------------|-------|
|   |            |        | Mean           | Best  |
|  | Mondrian   | Mean   | 21.76          | 16.79 |
|   |            | StdDev | 1.57           | 1.82  |
|  | Van Gogh   | Mean   | 16.51          | 12.45 |
|   |            | StdDev | 1.69           | 1.90  |
|  | Cable Ends | Mean   | 33.98          | 30.71 |
|   |            | StdDev | 2.11           | 2.70  |
|  | Flower     | Mean   | 14.87          | 10.76 |
|   |            | StdDev | 1.44           | 1.23  |

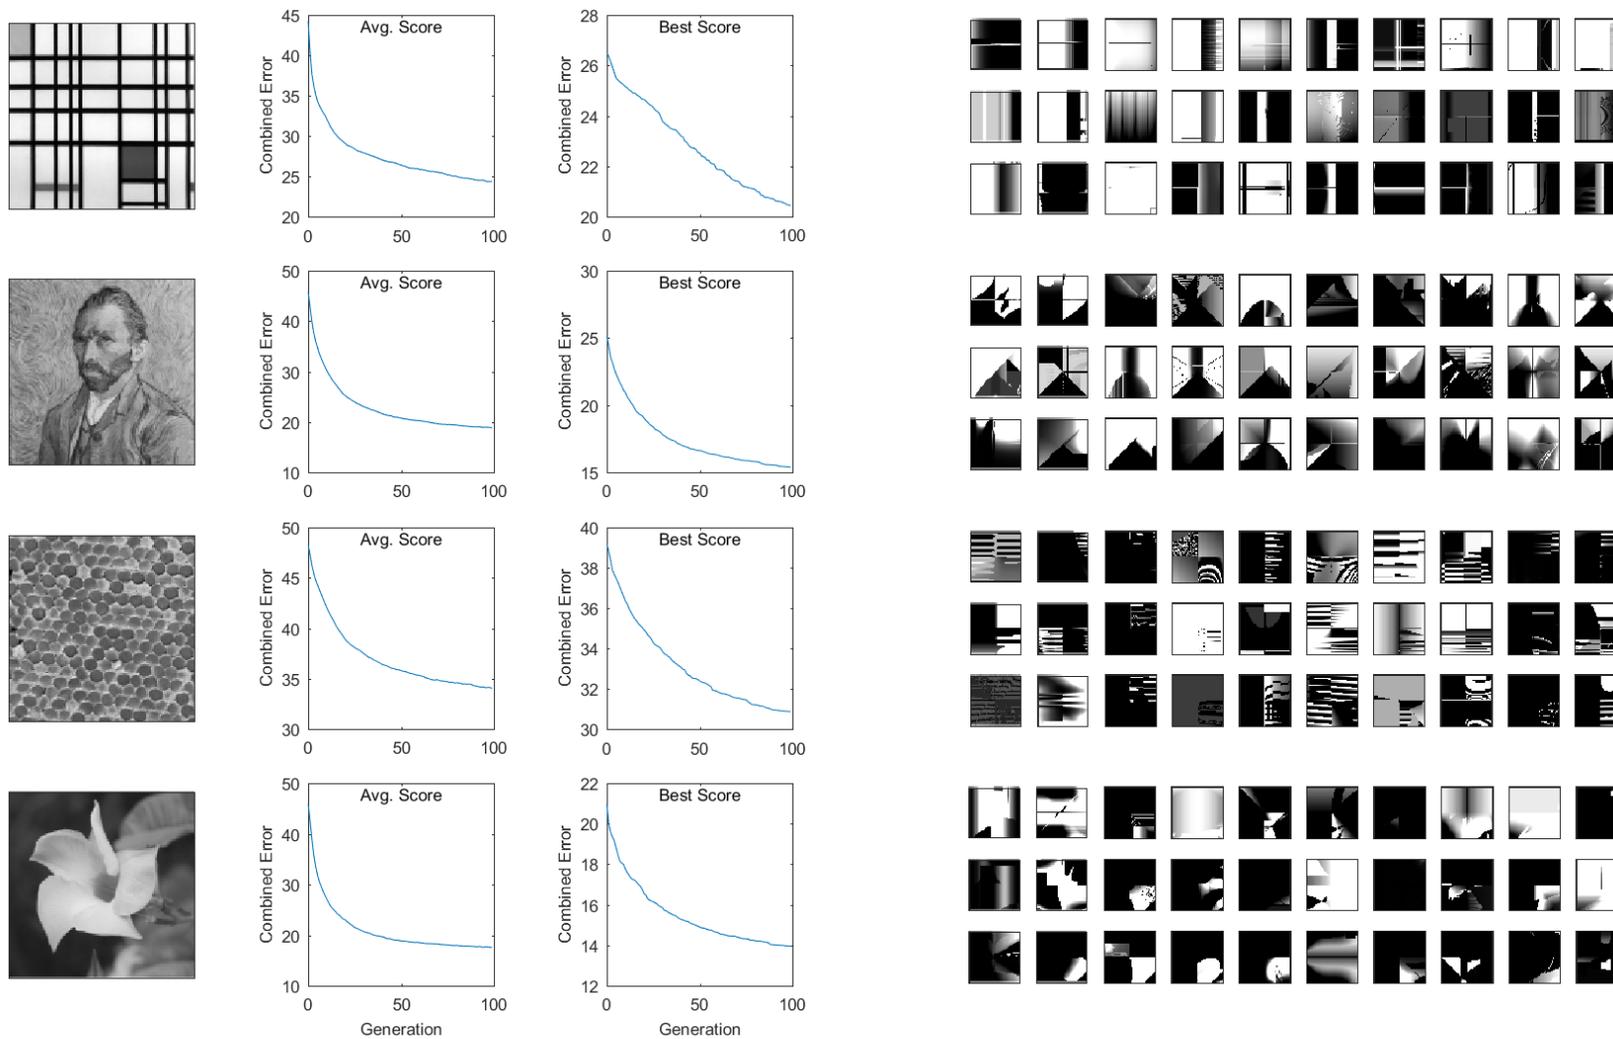
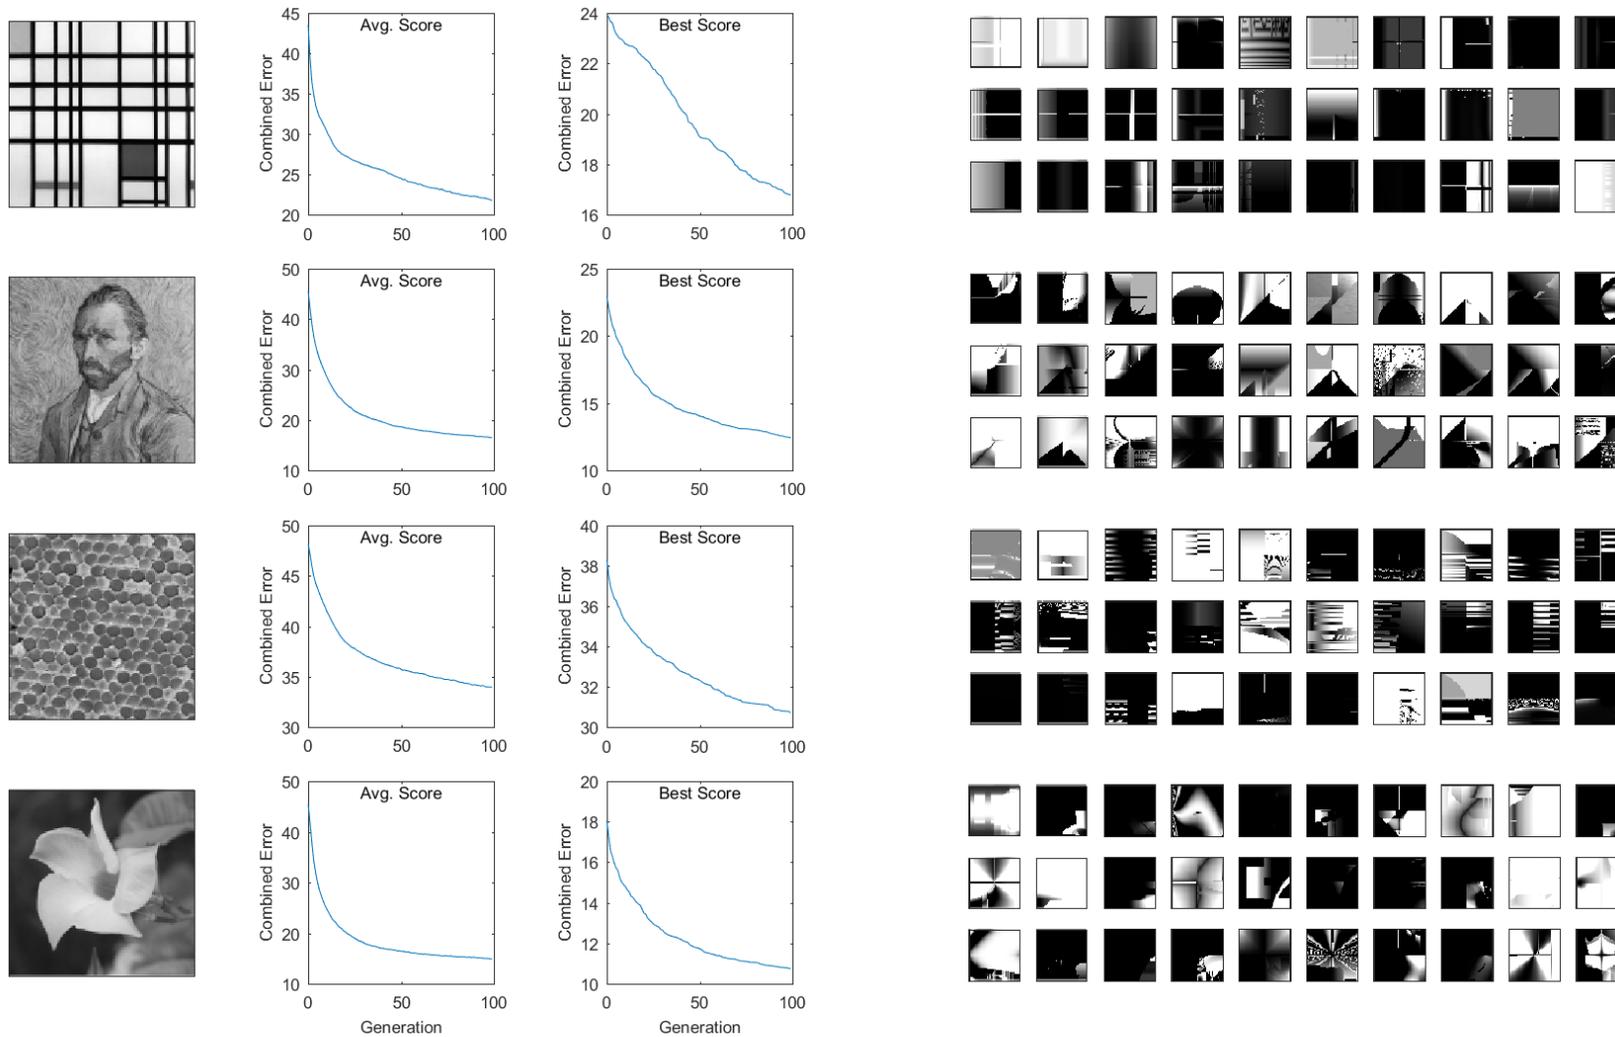
Figure 5.13: Experiment J3 Summary Charts & Examples;  $K = 50$ ; Linear Error

Figure 5.13: Experiment J3 Summary Charts & Examples;  $K = 50$ ; Square Error



scores to its J1 counterpart.

Comparing Table 5.9 to Table 5.7 (produced without phase penalties), we do not see a large increase in reported error. The best individuals of the run showed an average increase of only 1-2 points for normalized, squared phase error, though the population means increased with the expected 10-15 points (for  $K = 50$ , average of 0.2-0.3 from a max of 0.5 per position).

Originally performed due to an oversight in our system development, we include a variation of the current fitness scheme which used a linear, instead of square, phase error penalty. Using squared error is effective for exacerbating the effects of the more egregious coefficients in a set, and consequently adding pressure for evolutionary systems to identify them as problems. The consistently higher numeric scores from this experiment set raises some concerns over how much phase error is permitted due to weaker guidance away from —and inability to identify — the poorly aligned frequencies. However, there are some visual aspects and consistencies shown in these evolved candidates which might be preferable over the square error approach.

### 5.3.3 Power Coefficient Error & Objective Separation

Before proceeding to optimize our phase-adapted, top  $K$  mismatch approach, we return and consider some alternate adaptations which might make use of the full, non-quantized power coefficient errors. From our re-constructive analysis in the previous section, we see that the Mondrian target should still be capable of producing relatively recognizable reproductions when phase is ignored, but we do not see this being as prominent in actuality. We suspect, then, that there could still be some further improvements made to capture precise power information, perhaps also with the inclusion of phase.

We have seen that truncation of coefficients to the  $K$  most powerful has shown an improved effect on guidance during evolution. There are still concerns that even among the truncated set, the most powerful coefficients may be providing extreme pressures which overshadow the remaining coefficients. There are also related concerns that too large of a  $K$  value could dilute the guiding pressures of the measure, as we have seen previously. Thus, we will truncate coefficients to the choice of  $K = 50$ , which seemed suitable from our previous experiments, but consider schemes which use the full coefficient error.

Table 5.10: Experiment K1 Summary Table

|   | Target                | Agg.   | Combined Error |           |
|---|-----------------------|--------|----------------|-----------|
|   |                       |        | Mean           | Best      |
|  | Painting_Abstract_002 | Mean   | 47634110.79    | 499404.81 |
|   |                       | StdDev | 15649227.61    | 65703.39  |
|  | Painting_People_004   | Mean   | 25762310.89    | 2085.22   |
|   |                       | StdDev | 14424950.60    | 1064.29   |
|  | Pattern_Texture_012   | Mean   | 25789125.60    | 1426.70   |
|   |                       | StdDev | 17115758.74    | 132.07    |
|  | Photo_Objects_005     | Mean   | 23011363.10    | 29664.64  |
|   |                       | StdDev | 14479664.53    | 26257.33  |

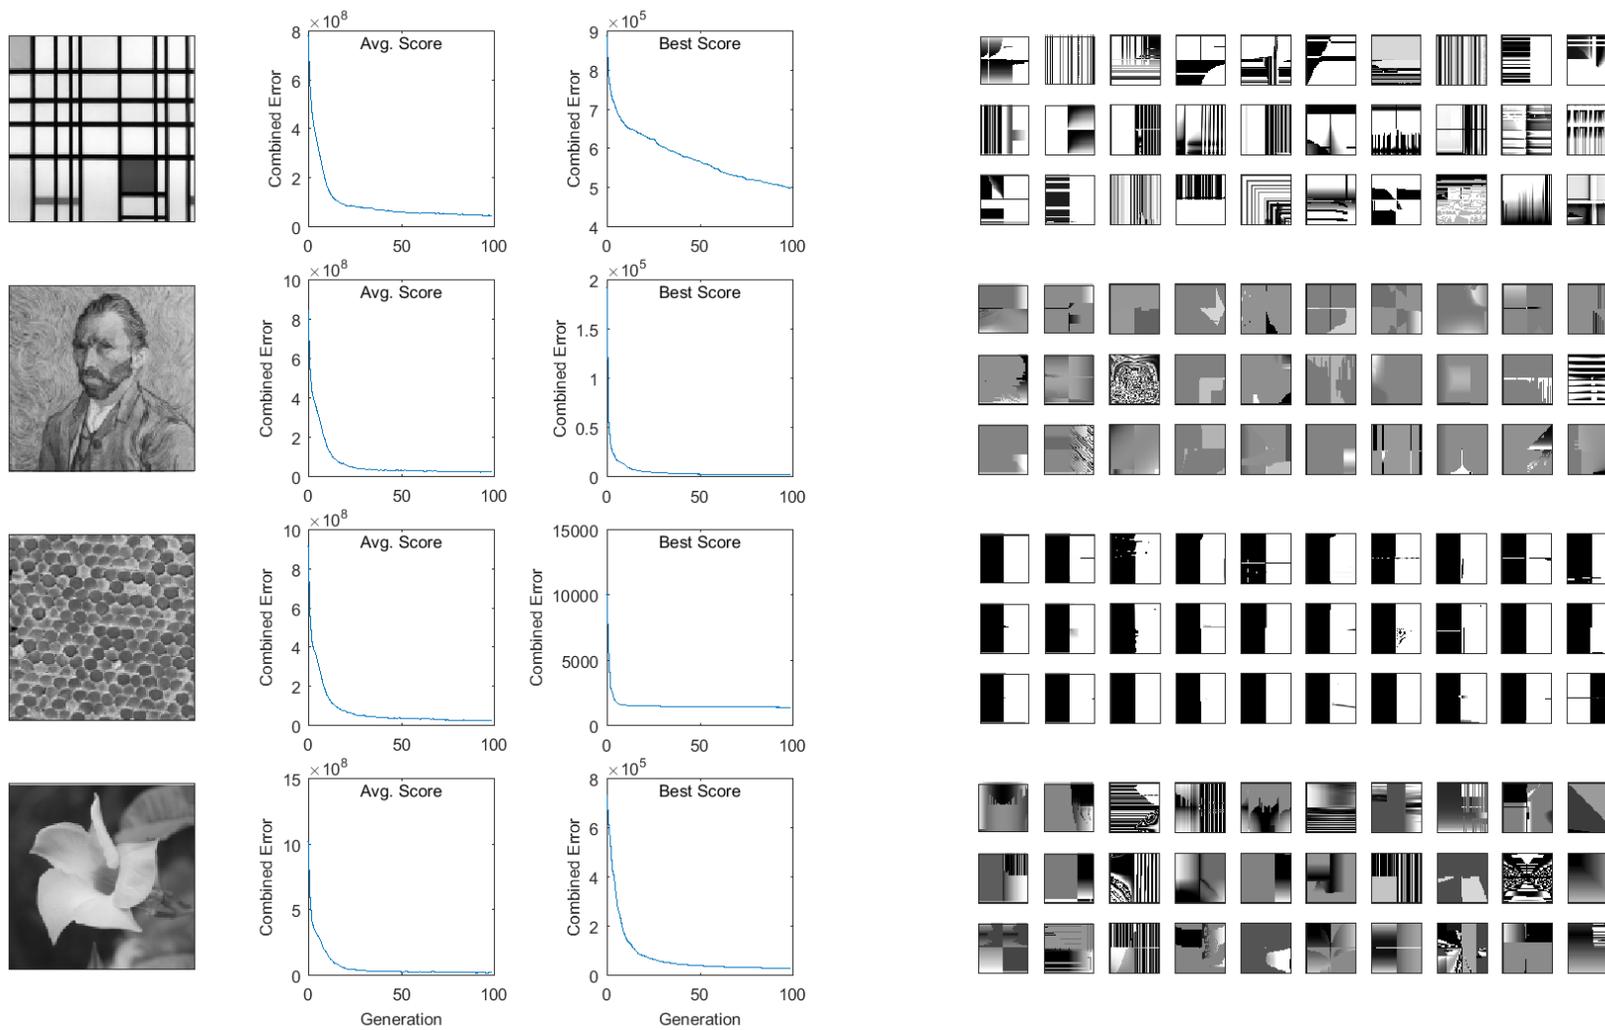
### Top Power & Phase Combined Coefficient Error

In our first return to power coefficient error, we will consider a single combined measure of power and phase across the top  $K = 50$  power coefficient positions. For each of the  $K$  positions, we will increment our error by the difference in power between target and candidate, and also by the squared, non-normalized (wrapped) phase angle error (ranged in  $[0, \pi^2]$ ). This should be similar to experiment J3, but using the full power coefficient values (instead of quantized to  $\{0, 1\}$ ), and a non-normalized phase angle error to hopefully prevent being completely overshadowed by the scale of the power error. We detail the approach in Equation 5.7 which uses the truncated set of coefficient positions for the target  $T$  and candidate  $C$  as ordered by power.

$$Error = \sum_{i=1}^K [ \Phi(T, T_i) - \Phi(C, T_i) ]^2 + [ \Delta( \theta(T, T_i), \theta(C, T_i) ) ]^2 \quad (5.7)$$

This borrows notation and functions from Equation 5.5, with  $\Phi(V, p)$  and  $\theta(V, p)$  returning the power and phase angle respectively of the coefficients (vectors) in set  $V$  corresponding to coordinate  $p$ .

Interestingly, the numerical results in Table 5.10 show similar properties to those found for experiment E4 (Table 5.5) in terms of relative error for the best individuals amongst targets. Visual contrast between Figure 5.14 and the previous Figure 5.5 from experiment E4 suggests that this approach may be producing more similar and novel results for some targets. Much like experiment E4, the Cable Ends target shows exceptional numerical performance despite little resemblance to the target. To some extent, the flower target maintains some levels of contrast as it had previously, and the Van Gogh target is still quite abstract. The Mondrian target, however, has much

Figure 5.14: Experiment K1 Summary Charts & Examples;  $K = 50$ 

more distinct lines in many of its evolved candidates. As the Mondrian target was the most improved, and its performance does not appear to heavily rely on phase, it would seem that there was some success in reducing the error produced through power coefficient mismatch, though it may not have benefited all targets.

### Top Power Coefficient Error, Separate Phase Error

With the results from K1 above, the addition of phase information and the truncation of the power coefficients to the top  $K$  positions appeared to provide us some benefits to the more power-dependent Mondrian target. However, the other targets which had previously shown some minor improvement when provided phase information did not appear to show notable changes from those produced in experiment E4. It may be advantageous to consider power and phase errors independently, both for the subsequent analysis of relative difficulty, and to potentially increase the evolutionary pressure for the targets showing difficulty in correcting phase.

Using the sum-of-ranks multi-objective fitness approach, we will isolate power error and phase error as two separate objectives for the subsequent experiment. Power error objective will be calculated from the  $K = 50$  truncated coefficients as the above experiment K1. The phase error objective will use the squared, normalized (wrapped) phase angle error from the same  $K = 50$  truncated coefficient positions. The formulae for the objective measures are seen in Equations 5.8 to 5.9.

$$Error_{power} = \sum_{i=1}^K [ \Phi(T, T_i) - \Phi(C, T_i) ]^2 \quad (5.8)$$

$$Error_{phase} = \sum_{i=1}^K [ \pi^{-1} \Delta( \theta(T, T_i), \theta(C, T_i) ) ]^2 \quad (5.9)$$

The ability to distinguish between power and phase errors in Table 5.11 provides us some possible insight into the difference of difficulty between targets. The poorly performing Cable End's target, despite showing exceptional power error scores, has a comparatively high phase score. This may help explain why we see candidates with poor visual quality, though this notion is offset somewhat by the target reconstruction in Appendix Figure A.1, which suggests that power should be the stronger requirement for this target. We can also observe that the Van Gogh target scores comparatively well across both objectives, despite many targets having very rough visual similarity. Between these two points, we are forced to reconfirm that difficulty

Figure 5.15: Experiment K2 Summary Charts & Examples;  $K = 50$

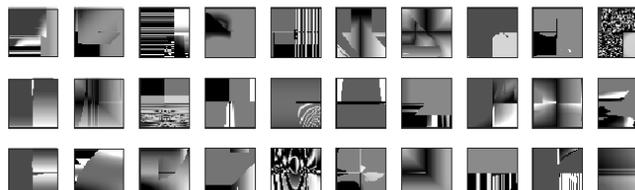
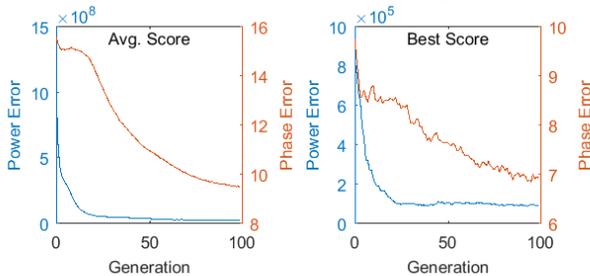
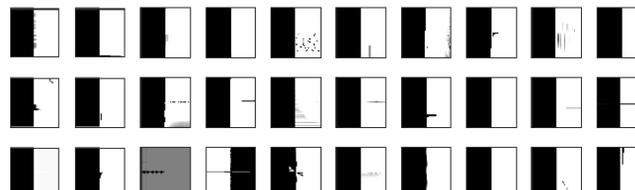
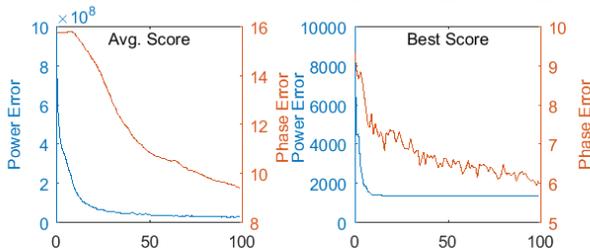
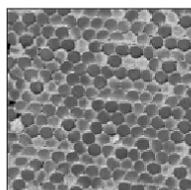
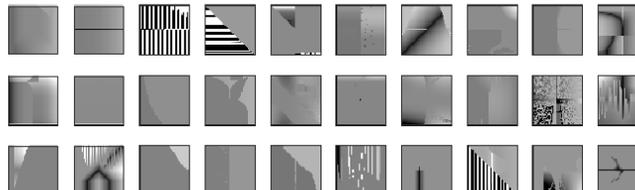
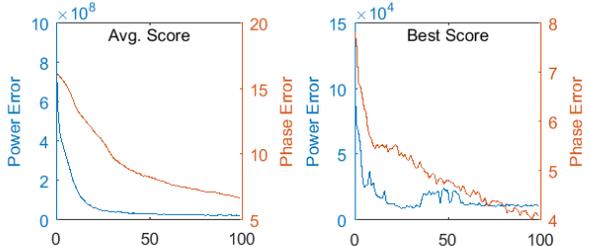
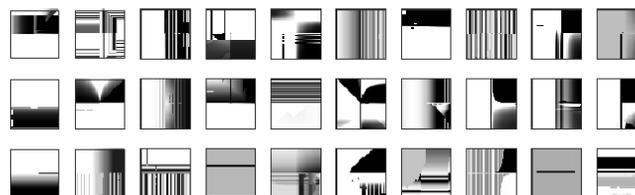
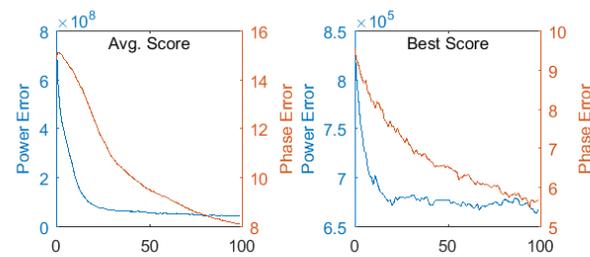
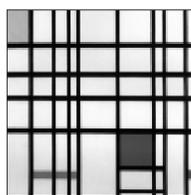


Table 5.11: Experiment K2 Summary Table

| Target  | Agg.   | Power       |           | Phase |      |
|---|--------|-------------|-----------|-------|------|
|   |        | Mean        | Best      | Mean  | Best |
|  Painting_Abstract_002 | Mean   | 47587277.68 | 668330.39 | 8.13  | 5.68 |
|   | StdDev | 23422058.74 | 73594.34  | 1.53  | 1.58 |
|  Painting_People_004   | Mean   | 22658547.22 | 10685.58  | 6.72  | 4.08 |
|   | StdDev | 11670497.09 | 10433.65  | 1.38  | 1.31 |
|  Pattern_Texture_012   | Mean   | 31659288.74 | 1397.18   | 9.40  | 6.09 |
|   | StdDev | 18502531.48 | 65.30     | 2.55  | 2.15 |
|  Photo_Objects_005     | Mean   | 23833558.68 | 91515.08  | 9.49  | 6.93 |
|   | StdDev | 16702114.90 | 97110.63  | 2.13  | 1.82 |

is highly target-dependent.

The visual performance shown in Figure 5.15 appears to have remained similar (if not slightly degraded with the Mondrian target) from the previous experiment. When we examine the performance plots for both the population mean and the best individual, we see some consistent behaviour across the two objectives. The power objective appears to improve and converge fairly quickly, where the phase objective is somewhat slower to stabilize. The phase objective is then either more difficult for our system to satisfy, the power objective is converging prematurely, or both. We will consider further refinements to the phase objective in the upcoming Section 6, but we will attempt one more experiment to adjust our handling of power coefficient error.

### Top Power Coefficient Error, Separate Phase Error, Separate 0th Error

Our last attempt to incorporate the non-quantized (though truncated) power coefficient error involves extracting the single most significant power coefficient and isolating it into its own additional fitness objective. The previous experiment left us with the consideration that there may still be heavy bias and premature convergence due to a small selection of key coefficients. The zero frequency, or  $0^{th}$  coefficient in the 2D DFT and power spectra relates to the overall offset of power required to recreate the overall intensity of the signal or image. In many cases with images, we will see this coefficient as the highest powered coefficient by a large margin (consider that image sample values are not negative, so any waveforms in the image will need to be positively offset). The  $0^{th}$  coefficient should be an important approximation of an images average intensity, and should be captured. We will nonetheless try isolating it to a separate objective so as to not unduly exclude the remaining power coefficients

Table 5.12: Experiment K3 Summary Table

(a)  $K = 50$ 

|   | Target                | Agg.   | 0th   |      | Power     |           | Phase |      |
|---|-----------------------|--------|-------|------|-----------|-----------|-------|------|
|   |                       |        | Mean  | Best | Mean      | Best      | Mean  | Best |
|  | Painting_Abstract_002 | Mean   | 11.65 | 0.00 | 421857.04 | 314059.28 | 11.03 | 7.65 |
|   |                       | StdDev | 5.15  | 0.00 | 80588.84  | 90657.25  | 3.03  | 1.86 |
|  | Painting_People_004   | Mean   | 4.60  | 0.00 | 669222.23 | 1909.08   | 9.43  | 5.60 |
|   |                       | StdDev | 4.51  | 0.00 | 455983.23 | 1081.65   | 2.90  | 1.74 |
|  | Pattern_Texture_012   | Mean   | 10.16 | 0.00 | 142984.81 | 941.34    | 10.02 | 6.58 |
|   |                       | StdDev | 7.33  | 0.00 | 107604.44 | 265.20    | 2.36  | 1.67 |
|  | Photo_Objects_005     | Mean   | 23.76 | 0.02 | 451154.15 | 14915.58  | 10.91 | 8.26 |
|   |                       | StdDev | 9.01  | 0.13 | 334762.94 | 3952.72   | 1.90  | 1.48 |

(b)  $K = 150$ 

|   | Target                | Agg.   | 0th   |      | Power     |           | Phase |       |
|---|-----------------------|--------|-------|------|-----------|-----------|-------|-------|
|   |                       |        | Mean  | Best | Mean      | Best      | Mean  | Best  |
|    | Painting_Abstract_002 | Mean   | 10.14 | 0.00 | 835566.64 | 476798.51 | 40.11 | 33.20 |
|   |                       | StdDev | 5.14  | 0.00 | 235074.43 | 69166.89  | 5.32  | 3.66  |
|    | Painting_People_004   | Mean   | 4.12  | 0.00 | 631053.51 | 3036.02   | 36.54 | 28.43 |
|   |                       | StdDev | 4.24  | 0.00 | 396339.35 | 1389.43   | 7.16  | 4.97  |
|   | Pattern_Texture_012   | Mean   | 8.72  | 0.00 | 197305.03 | 1123.29   | 39.06 | 32.62 |
|   |                       | StdDev | 7.57  | 0.00 | 184609.25 | 284.71    | 3.94  | 3.04  |
|  | Photo_Objects_005     | Mean   | 17.72 | 0.00 | 224754.35 | 14258.83  | 36.09 | 31.58 |
|   |                       | StdDev | 5.84  | 0.00 | 104786.48 | 3609.90   | 3.26  | 2.95  |

from providing evolutionary guidance. The existing two sum-of-ranks fitness objectives from the previous experiment, power error and phase error, will be maintained in addition to this new objective. These three objectives are detailed in Equations 5.10 to 5.12.

$$Error_{0th} = | \Phi(T, T_1) - \Phi(C, T_1) | \quad (5.10)$$

$$Error_{power} = \sum_{i=2}^K [ \Phi(T, T_i) - \Phi(C, T_i) ]^2 \quad (5.11)$$

$$Error_{phase} = \sum_{i=1}^K [ \pi^{-1} \Delta( \theta(T, T_i), \theta(C, T_i) ) ]^2 \quad (5.12)$$

The greatest power coefficient is expected to correspond to the  $0^{th}$ -frequency, and thus we have position for  $T_1$  lie at the center of the power coefficients (  $\{ \frac{n}{2}, \frac{n}{2} \}$  ).

When observing the difference of performance between the previous experiment

Figure 5.16: Experiment K3 Summary Charts & Examples;  $K = 50$ 

While plot curves, labels, and axes are usually associated by colour, there is a notable exception in this figure (and the following variant for  $K = 150$ ), to accommodate the scale of the additional  $0^{th}$  coefficient objective. The dashed red lines correspond to error between the  $0^{th}$  power coefficient, but despite being a power-related objective, it will be scaled along the phase axes to better display the plot behaviour.

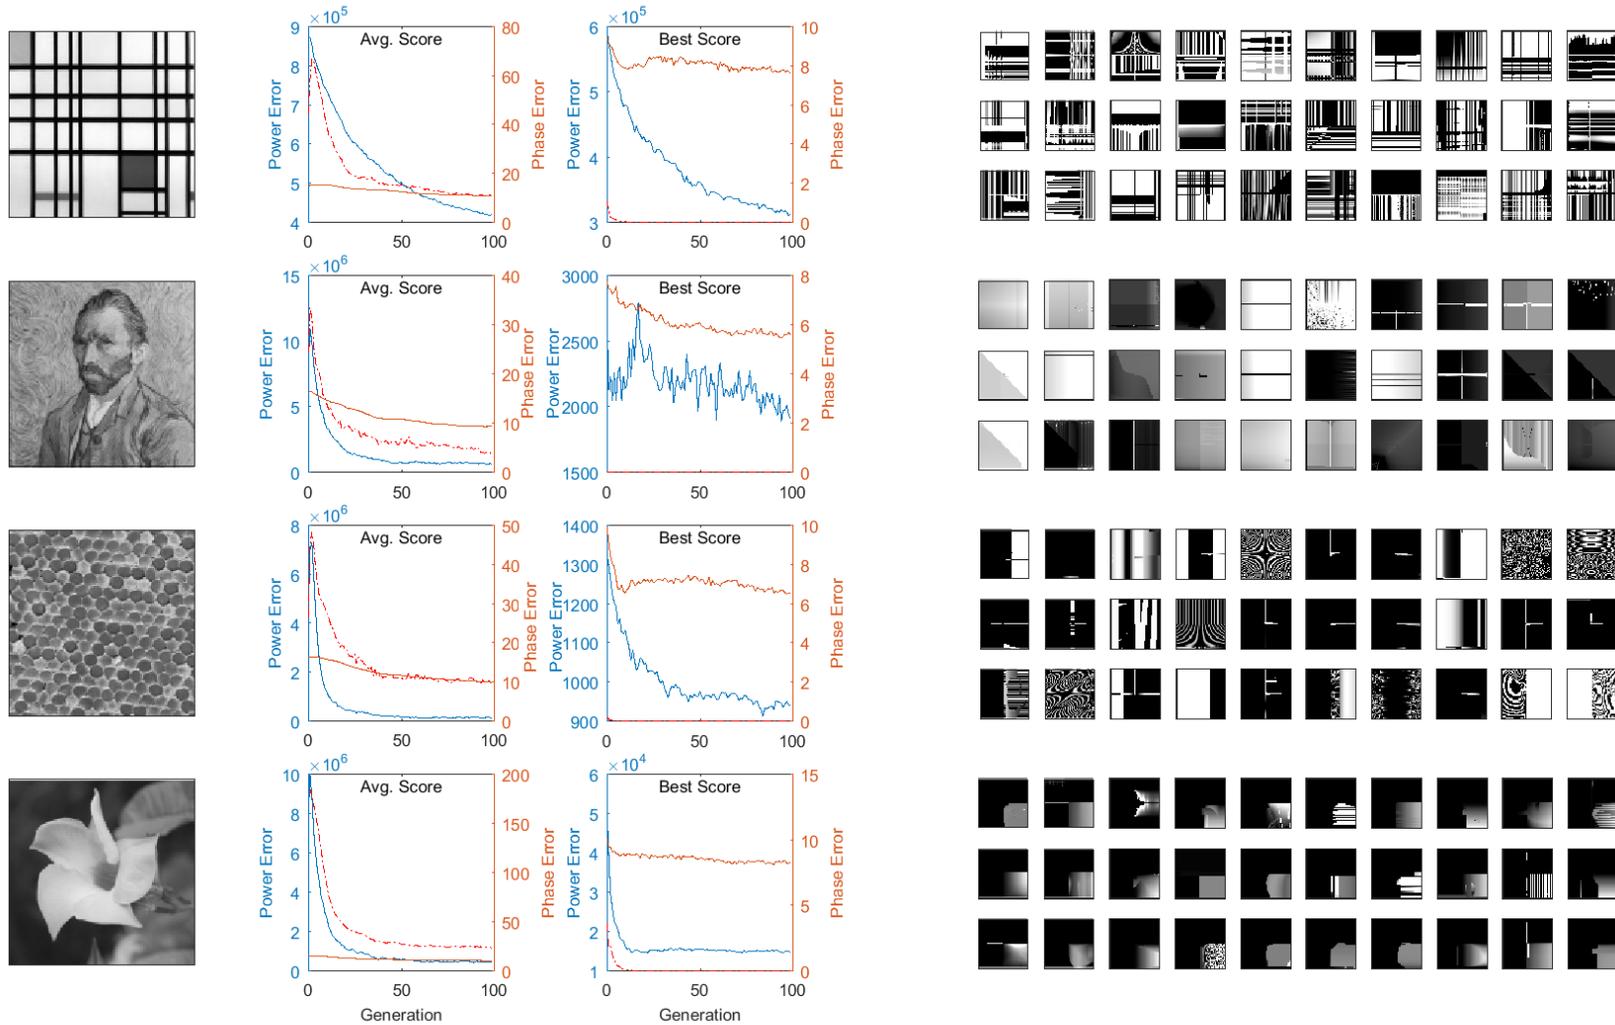
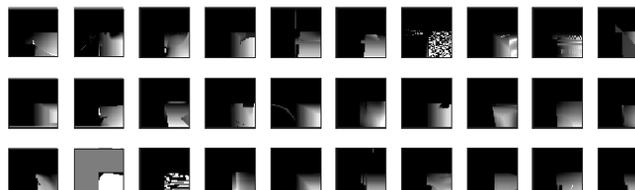
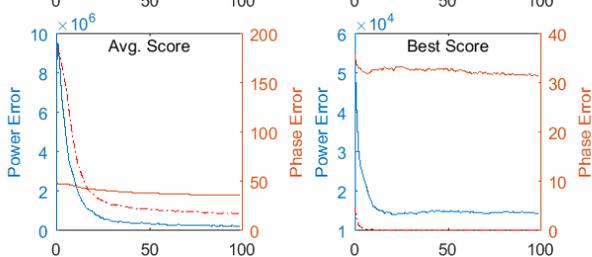
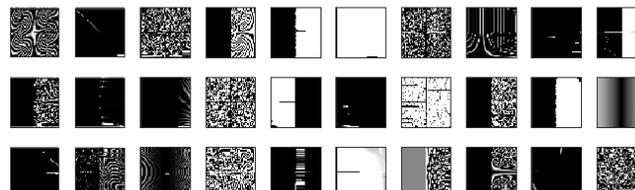
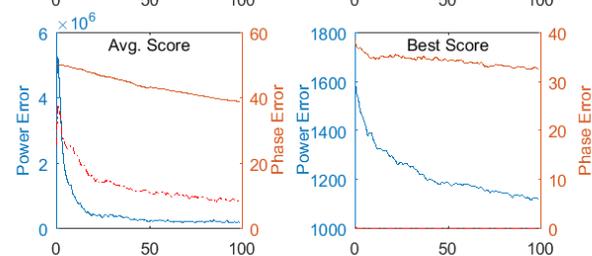
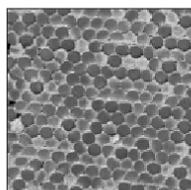
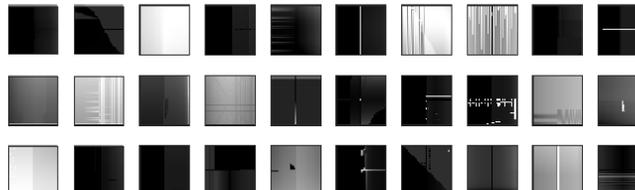
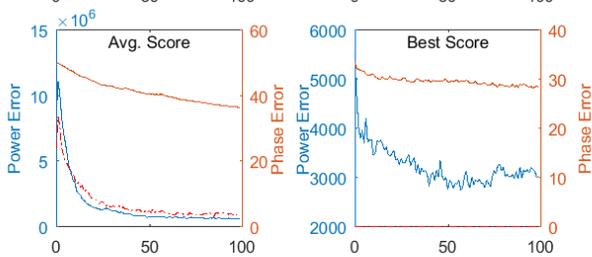
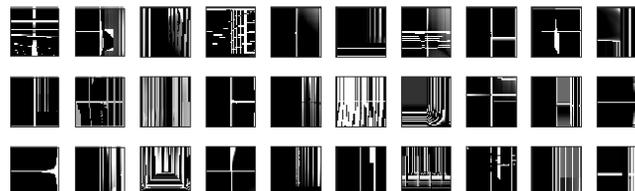
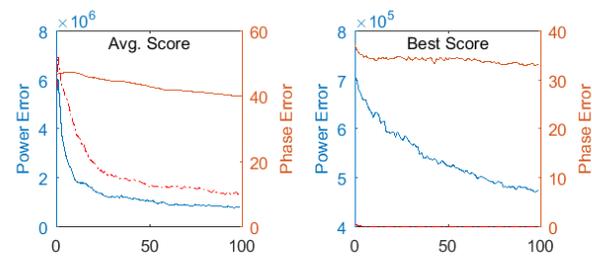
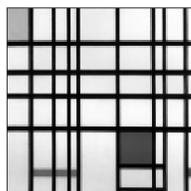


Figure 5.16: Experiment K3 Summary Charts & Examples;  $K = 150$



and the results for  $K = 50$  in Table 5.12, there is a substantial improvement to power scores. The 0<sup>th</sup> coefficient objective has been consistently minimized to an error of  $< 0.01$  for most targets, while scores for the power objective have substantially decreased (by factors between 1.5 to 5 times). Phase has worsened only slightly, which was not altogether unexpected considering the fixed generation count, and additional objective to be optimized. It would seem that by isolating the 0<sup>th</sup> objective, we have accomplished the goal of allowing the remaining coefficients to show greater effect.

Examining the performance plots in Figure 5.16, we do largely see similar behaviour to the previous experiment, so it is not fully clear if we have resolved our concerns of premature convergence. However, many of the targets do show a noisier refinement of their best individuals. While we have seen some competition between objectives in the previous experiment, this may be a promising signal that the power objective has additional possibilities to explore in its guidance of the candidate images.

Figure 5.16 also shows some long-awaited visual improvements. While the Van Gogh target remains abstract, the Cable Ends target begins to show some higher-frequency patterns, and the flower target seems to consistently capture and evolve either the lightened corner, or out-of-phase petal cluster. Most notable is the improvement to the Mondrian target, which now solidly displays similar features to the target: crossing lines and intensity blocks.

Results with  $K = 150$  had been omitted for the previous two experiments as there were no notable visual changes observed. With the current fitness measure, the change in  $K$  leads to mixed results. The Van Gogh and flower targets see little change, and the Mondrian target appears to give less ideal visual features. The ever-difficult Cable Ends target, with the increased  $K$ , shows an improved consistency in its higher-frequency, seemingly noisy patterns. This is curious, as the Cable Ends targets can provide reasonable reconstructions at only  $K = 10$ . A current suspicion is that some of the higher-powered coefficients for this target are difficult to be guided toward without intermediate frequencies as targets. This again signals the difference in difficulty between certain targets.

# Chapter 6

## Compositional Refinement

Previous experiments have provided us with numerous fitness approaches to build upon. Our exploration with power coefficient error schemes has resulted in an approach which may be suitable for evolving basic targets which do not require high detail in phase information. However, many images will not be satisfied by these constraints, and we are still somewhat lacking in our ability to guide coefficient phase angle. In this section, we will adjust our fitness schemes to obtain a measure which is capable of guiding and producing candidate images with phase-dependent, compositional attributes.

Table 6.1: Overview of Section 6 Experimental Variations

| Label | Measure | K  | Language | Tables | Figures | Notes  |
|-------|---------|----|----------|--------|---------|--|
| J3    | J3      | 50 | Base     | —      | 6.2     | Review of J3 (current best quantized measure) with new compositional targets     |
|       | J3      | 10 | Base     | —      | 6.4     |  |
| K3    | K3      | 50 | Base     | —      | 6.3     | Review of K3 (current best non-quantized measure) with new compositional targets |
| P1    | P1      | 50 | Base     | 6.2    | 6.5     | New, general-purpose measure   |
|       | P1      | †  | Polar    | —      | 6.6     | K value and language adjustments for survey preparation                          |

† Choice of K in these experiments varies per target. Selections are outlined in Table 6.3, and reiterated within the tables of this section.

We introduce new targets to assist with our exploration, upon which we will first review our current top-contending measures. We then use our findings to produce a new measure, which we recommend for its fair performance across all targets. The experiments of this section are outlined in Table 6.1.

## 6.1 Compositional Targets

The initial target set was chosen to express a varied selection of our identified image genres and styles. Unfortunately, by this point we are still largely unable to capture any of the finer style traits from the images in the initial target set. Contrarily, we find that attempting to quantify and guide candidates with similar coarse composition to be more tractable. Where we expect finer, micro-scale style properties to be captured by the higher frequency power components, we see macro-scale compositional properties to be more easily expressed through the dominating low frequency components. We have also begun to suspect that our choice of a simple GP language may be adding to the difficulty in reproducing the finer target image traits of the comparatively complex target images.

At this point, we refocus our efforts on the aspect of spatial composition similarity. We consider a new set of target images which are visually simpler than our initial set, but provides additional variation in basic but common image compositions, and varied demands of phase precision. These new target images include compositions frequently borrowed for evolutionary art, and we hope that they will be capable of giving us more insights when refining our fitness measures to his end.

Figure 6.1 outlines the new compositional targets used for the refinement of our phase related measures. Power spectra coefficient displays and reconstructions of the new targets can be found in Appendix A, Figure A.1.

## 6.2 Revisiting Previous Measures

Before attempting to refine some of our previous measures, we first familiarize with how our new target images perform with the existing fitness measures.

When we examine the performance of our previous measures on the new target set (Table 5.1 can briefly recap our explored measures), we see some generally mixed results. Across both J3 (Figure 6.2) and K3 (Figure 6.3), targets Composition\_02 and Composition\_03 (black and white split) very consistently produce near-perfect reproductions in the initial or early generations. Conversely, targets Composition\_07

Figure 6.1: Compositional Target Set

A number of the targets are fundamentally similar, but hold variations which we believe may provide difficulty for our measures. Composition\_03 differs from Composition\_02 by adding a rotational aspect, which is further varied in Composition\_04 by changing the proportion of high and low pixel intensity. Composition\_06 builds upon Composition\_05 by increasing the size and number of stripes, and in doing so, transitions from a regular (repeating) texture to an irregular texture. Composition\_08 adds translation to every odd row from Composition\_07 (incidentally creating a diagonal repetition), and Composition\_12 adjusts the frequency of the overlapping Gabors from Composition\_11 to create different interference effects. Composition\_01 and Composition\_09 are included as examples of basic but common patterns that may be found in abstract and minimalist artwork.

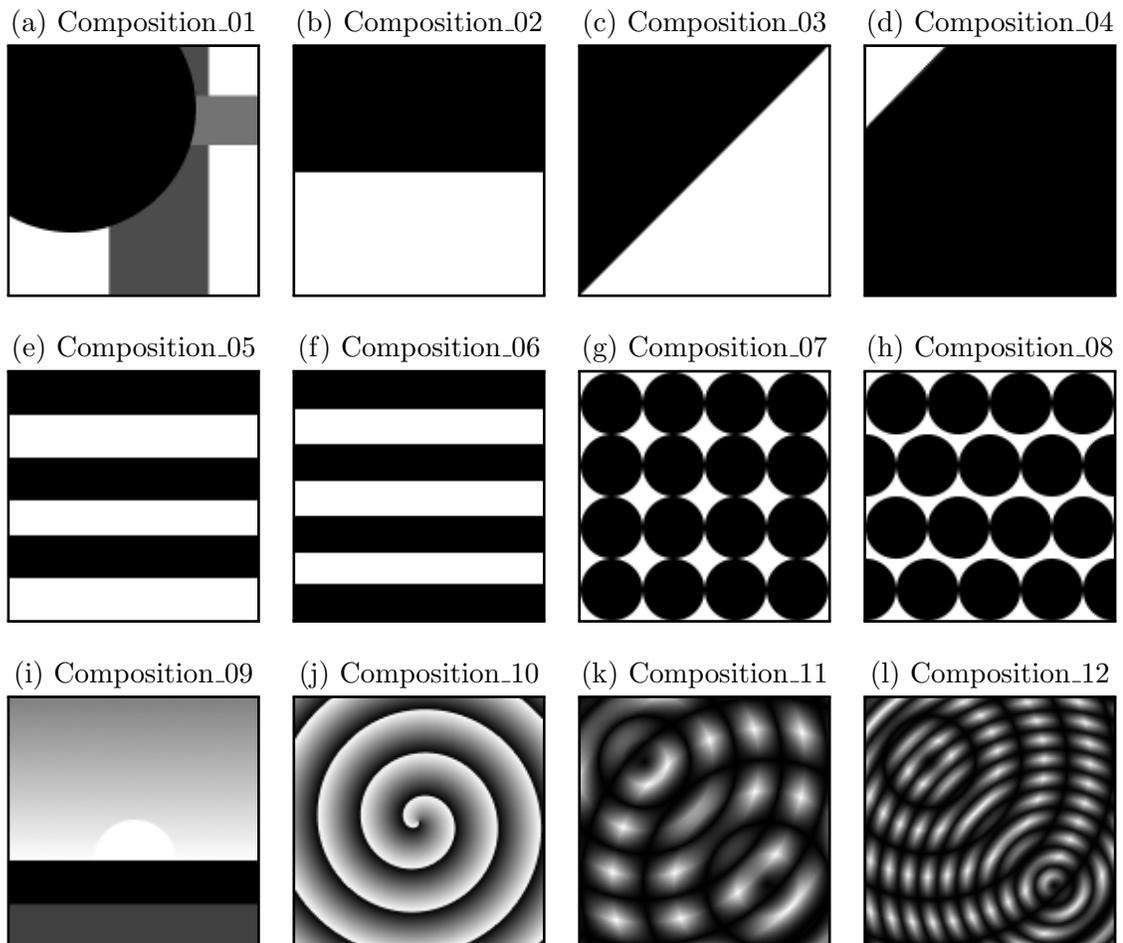


Figure 6.2: Experiment J3 Compositional Summary Charts & Examples;  $K = 50$

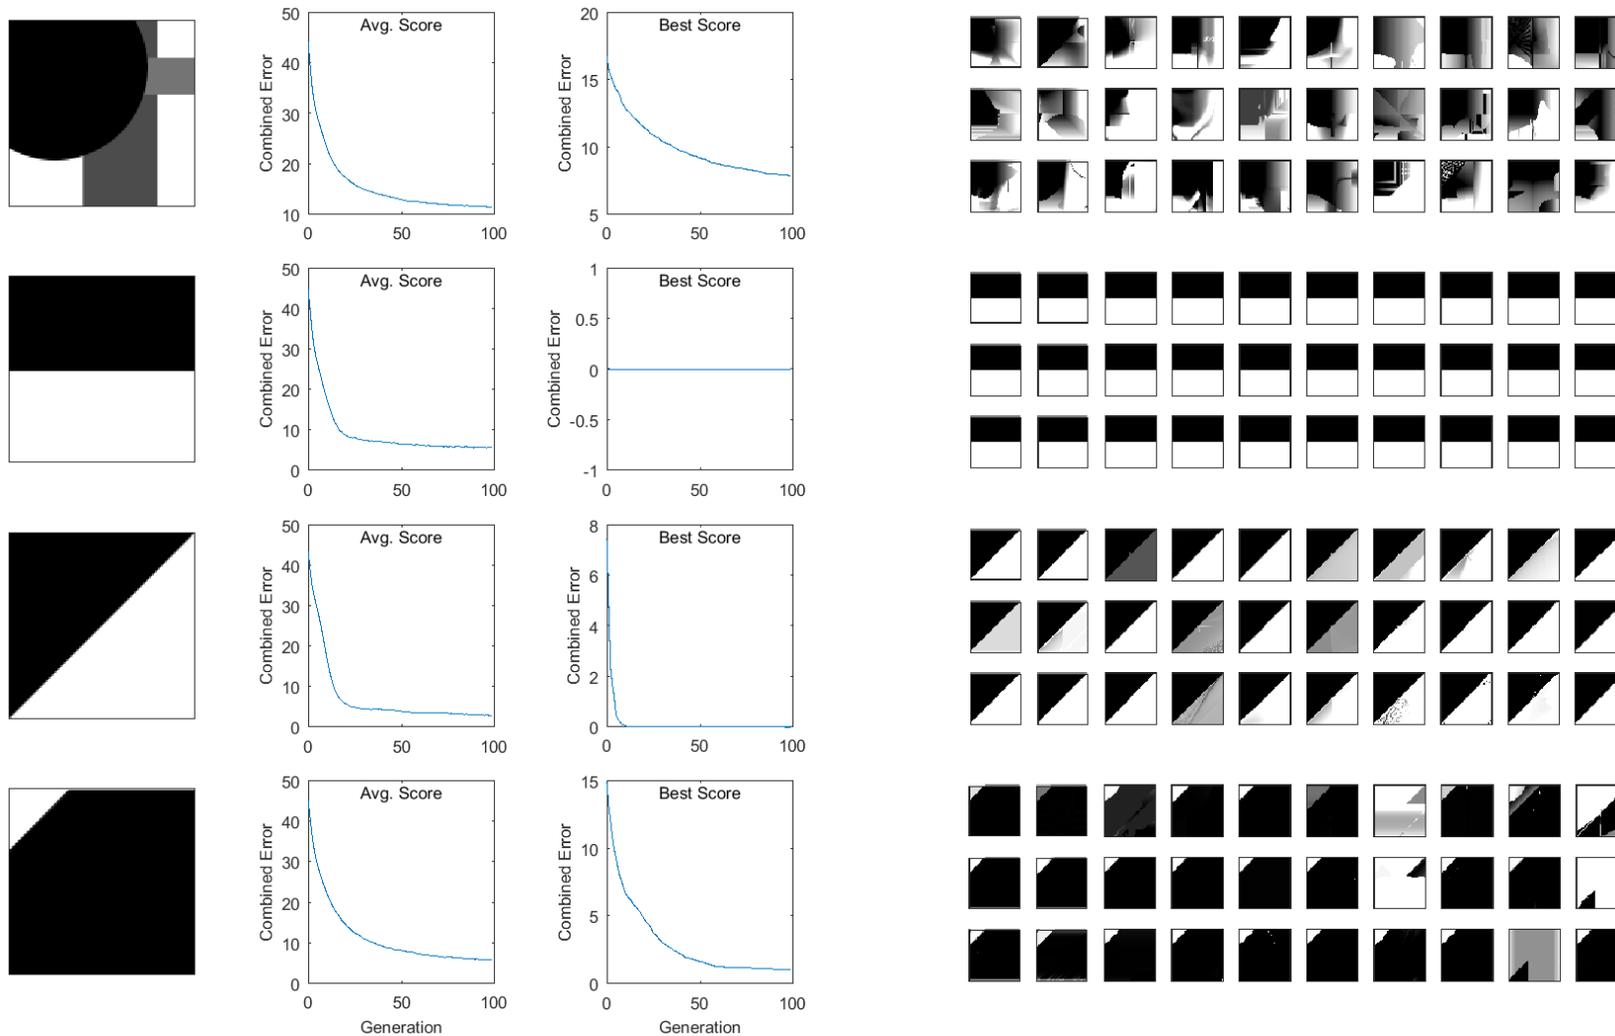


Figure 6.2: Experiment J3 Compositional Summary Charts & Examples;  $K = 50$

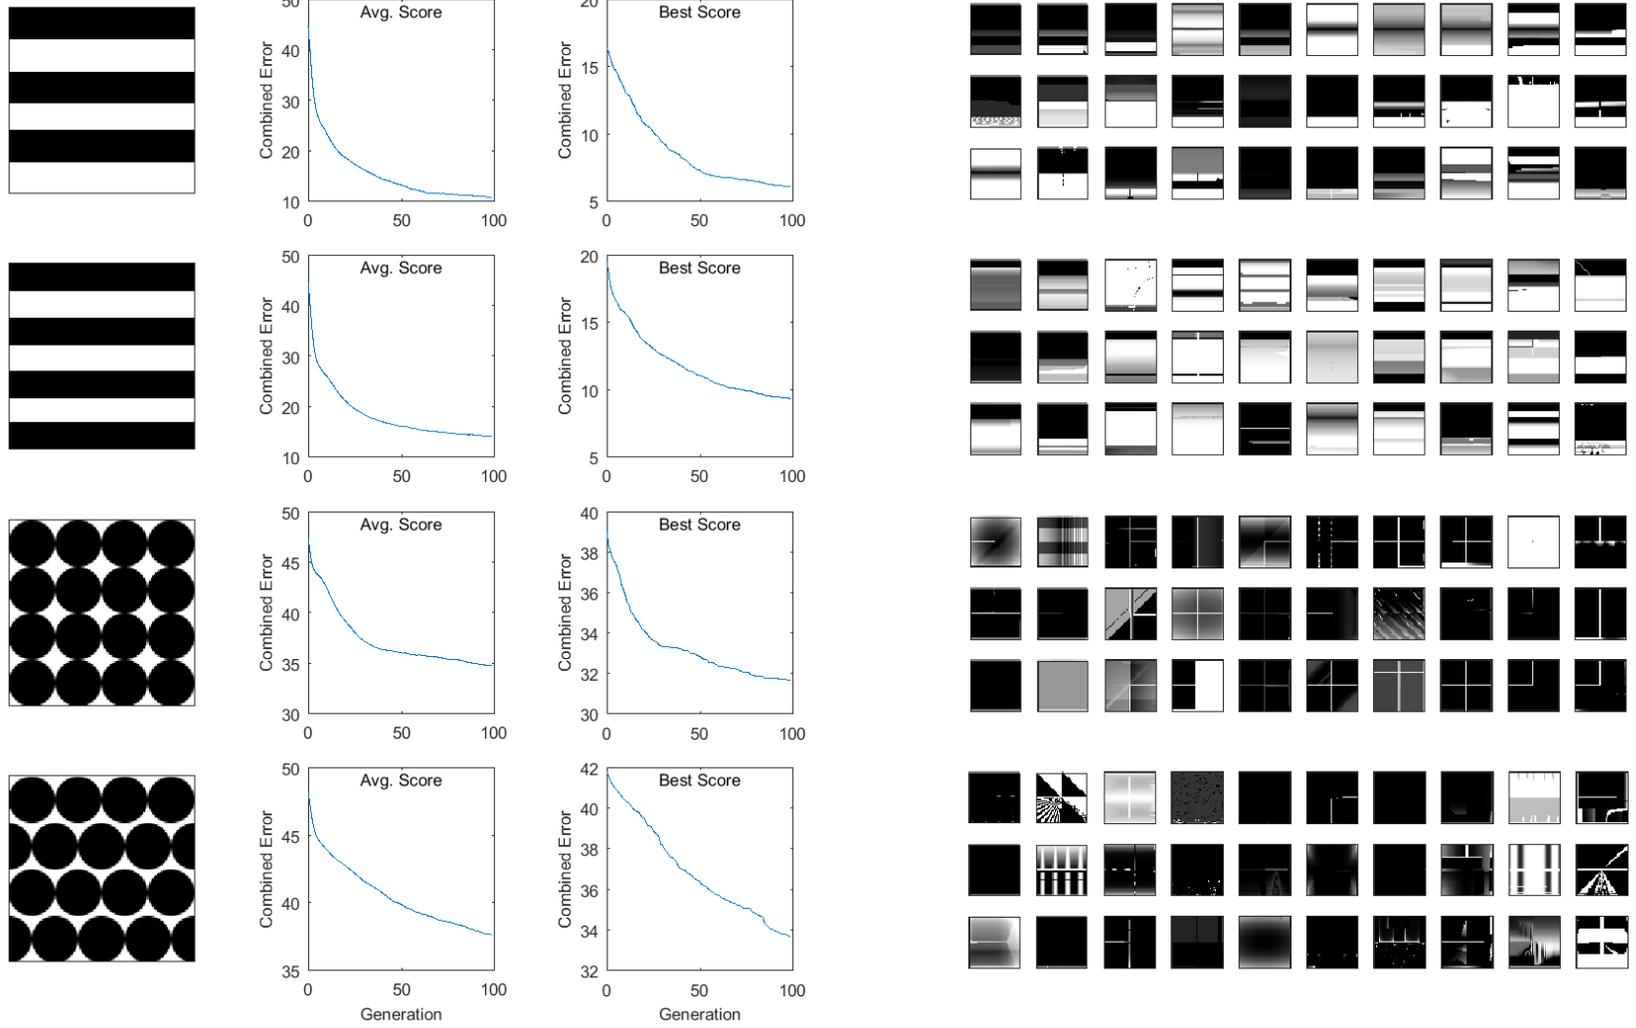


Figure 6.2: Experiment J3 Compositional Summary Charts & Examples;  $K = 50$

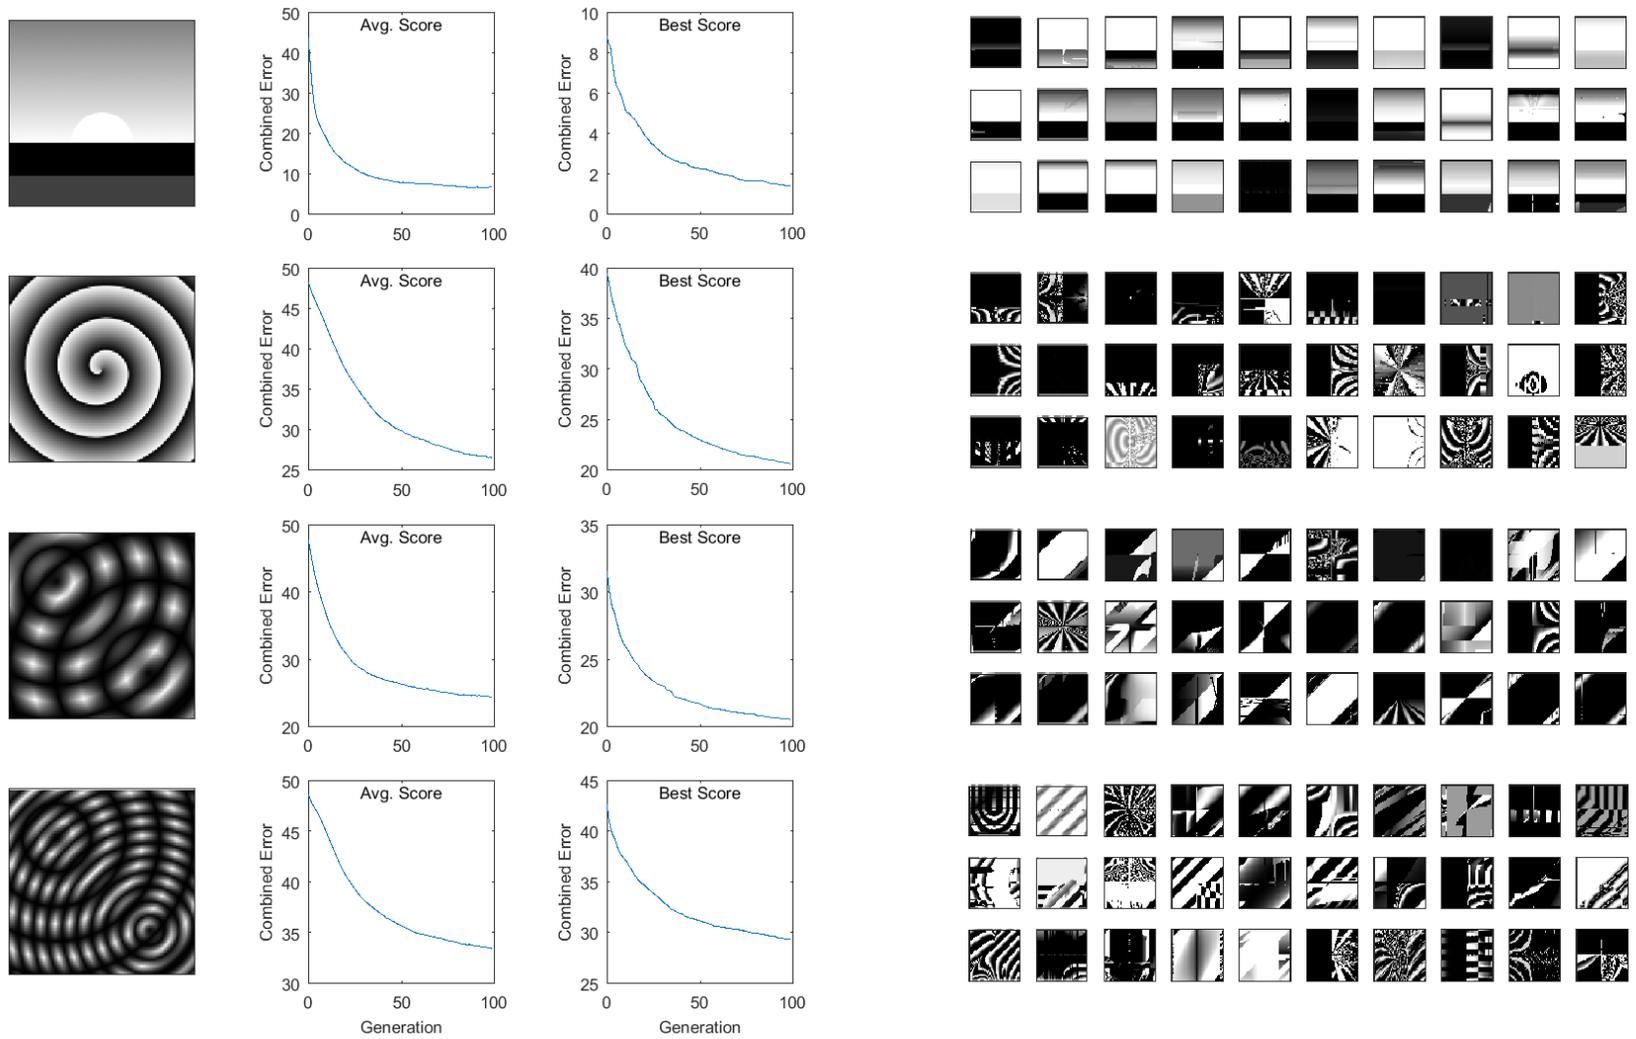


Figure 6.3: Experiment K3 Compositional Summary Charts & Examples;  $K = 50$

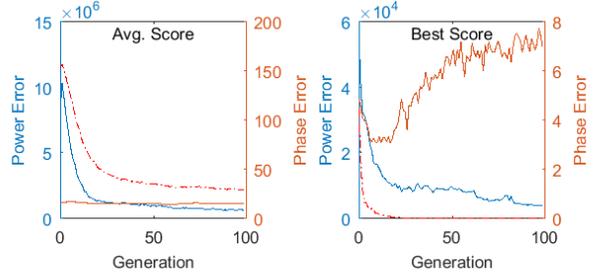
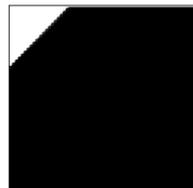
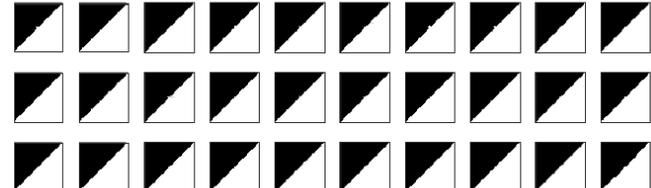
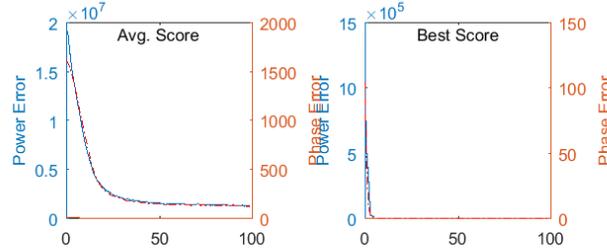
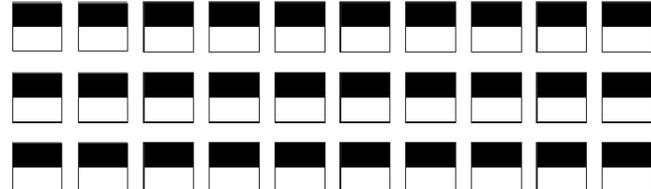
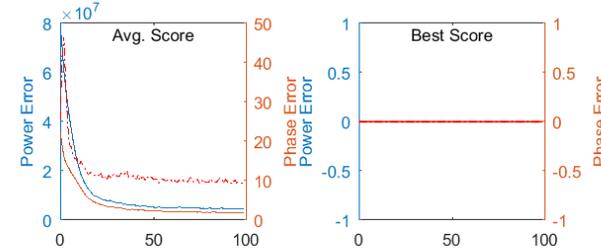
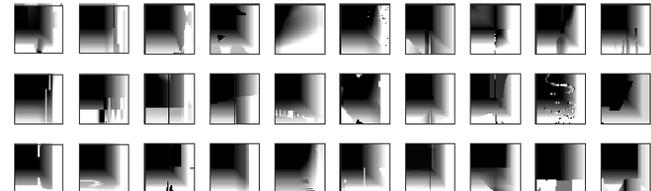
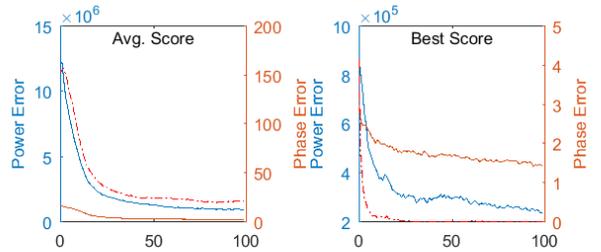
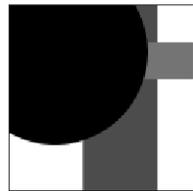


Figure 6.3: Experiment K3 Compositional Summary Charts & Examples;  $K = 50$

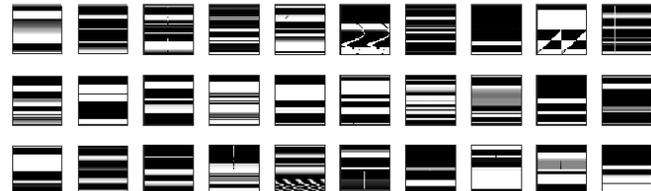
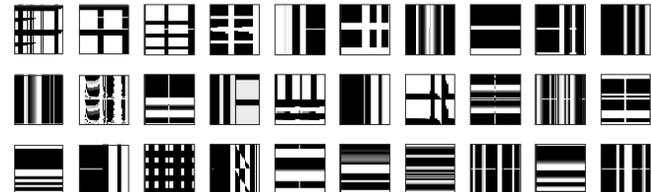
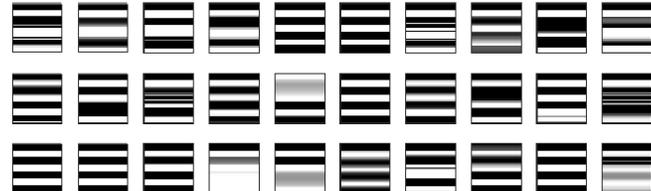
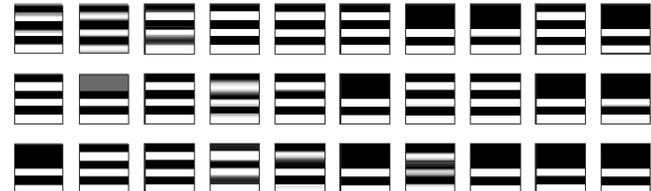
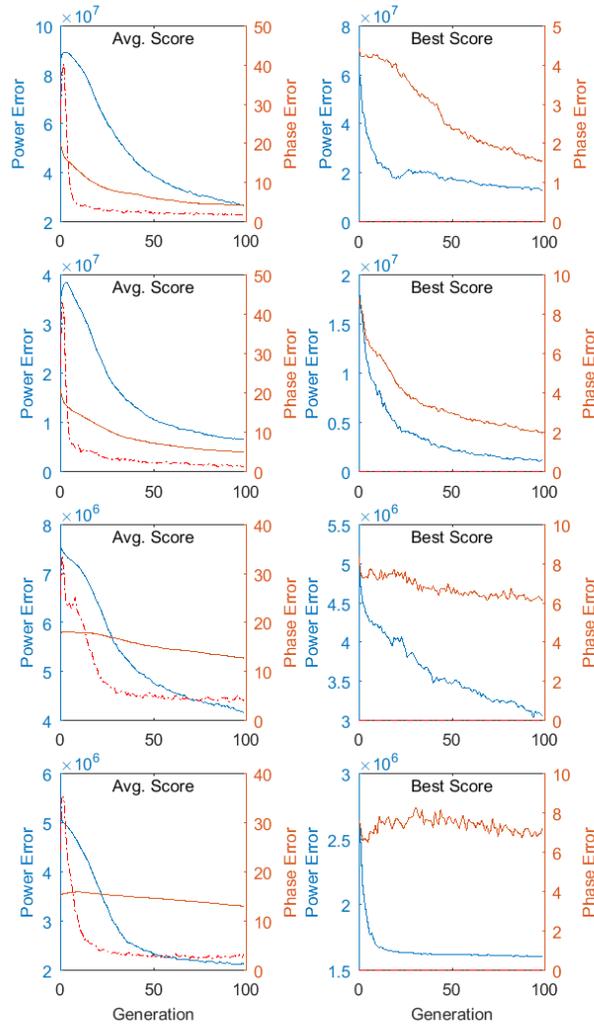
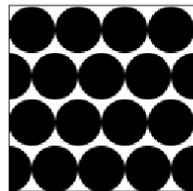
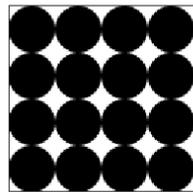
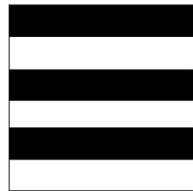
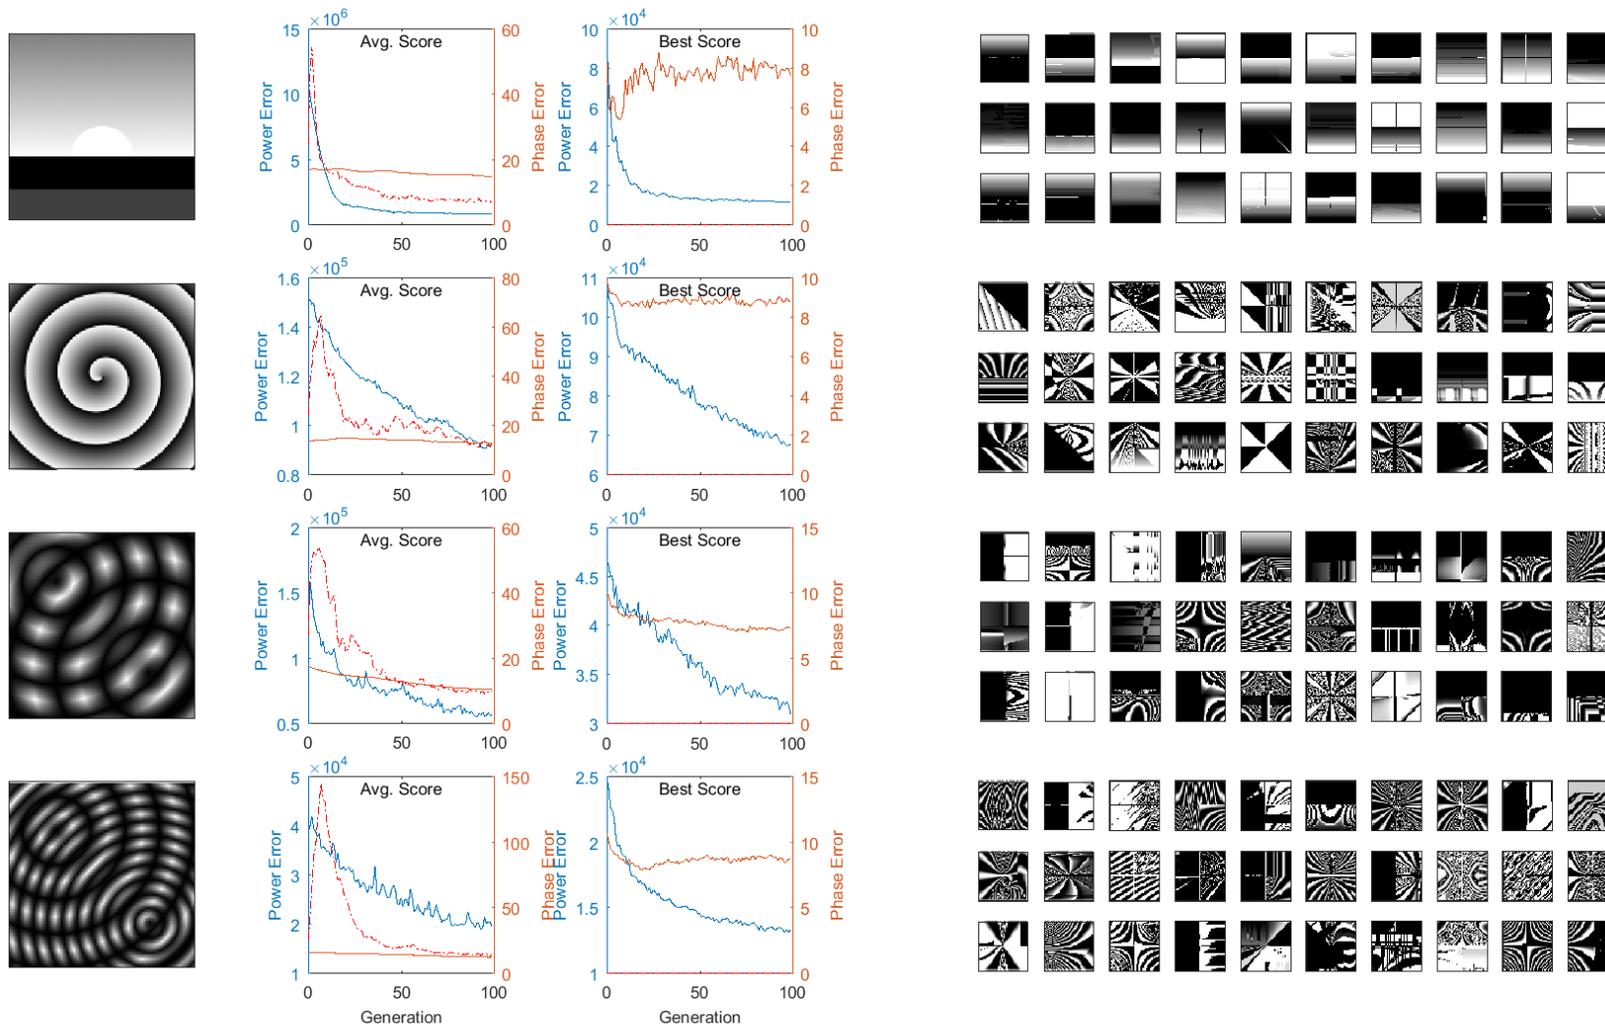


Figure 6.3: Experiment K3 Compositional Summary Charts & Examples;  $K = 50$ 

and Composition\_08 (the circle grids) remain difficult to reproduce throughout the remainder of this chapter. The remaining targets show various degrees of success, reinforcing that difficulty of the measures is highly target dependent, and affording us space to improve with these targets. Composition\_01 and Composition\_09, the minimalist art compositions, give vague but reasonable and interesting candidate recreations with both of our existing measures. Neither the various spiral/Gabor targets (Composition\_10-12) nor the circular grid targets (Composition\_07-08) show any particularly strong circular aspects, though we do see similar contrast frequencies and chaotic behaviour with the former, and consistent grid behaviours with the later.

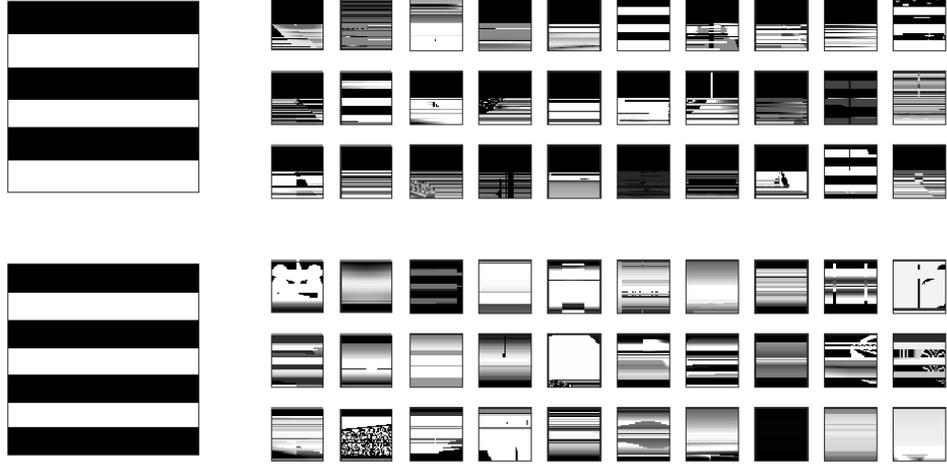
Some interesting divergences appear between the two measures with a number of the new compositional targets. With the striped targets, we see more faithful reproductions using the K3 measure, which previously improved the striped Mondrian target and other less phase-dependent targets. With Composition\_04, we see fairly consistent reproductions using the measure from J3, which quantized the power error and used a single mixed objective with phase. In contrast, when using the measure from K3, we see that a triangle similar to that of the target is produced, but placed offset from the corner of the image – out of phase. When further evaluating the plotted scores for this target with K3, we see that between the three fitness objective, the best individual’s phase error objective was sacrificed and steadily increases as generations progress. Many of the other targets show similarly high phase error values for their best individuals.

We have also explored these targets and fitness measures with lower values chosen for  $K$ . Figures B.1 and B.1 in Appendix B show results at  $K = 10$ . Many of these targets performed very well truncating the power coefficients to  $K = 10$ , and notably, the stripes in Composition\_05 and Composition\_06 showed improvements with the J3 measure (Figure 6.4). With the reduced  $K$  and proportional reduction to error scores, identical concerns were still noted regarding the sacrificed phase error objective.

We suspect that to optimize a more generalized, all-purpose measure, we will need to increase the pressure from the phase objective, through objective weighting, reduction of objectives, relaxing power error through quantization, phase error scaling schemes, or some combination of the previous approaches.

### 6.3 Adjusting Phase Error

After a few revisions inspired by the findings of our previous experiments, we have found a measure which appears the most promising when considered across all of our

Figure 6.4: Experiment J3 Compositional Examples;  $K = 10$ , Composition\_04-05

targets. Returning to J3 – which until now showed the most promise regarding phase accuracy – we separate the phase error component to its own sum-of-ranks fitness objective, and apply a scaling factor to the phase to prioritize the more powerful components. This is more formally defined in Equations 6.1 and 6.2, and will be used for our next experiment, P1. While the non-quantized approach performed well with experiment K3 for the less phase-dependent targets, it was at the cost of degraded performance of the other targets. Re-examining the quantized J3 measure with our new compositional targets, we found reasonable performance across the entire target set so long as  $K$  was tuned appropriately per target. From our observations of phase in experiment K3, we saw phase was often sacrificed in lieu of power due to what we presume was a higher inherent difficulty with that aspect during evolution. To try and accommodate this phase objective, we add a linear scaling to each part of the phase error sum, such that errors seen at the more powerful coefficient positions were correspondingly more meaningful. If the phase measure cannot be completely fulfilled during evolution, we hope we can now prioritize the offsets of the more visually prominent (powerful) frequencies.

$$Error_{power} = \sum_{i=1}^K \begin{cases} 0.0 & , T_i \in C \\ 1.0 & , T_i \notin C \end{cases} \quad (6.1)$$

$$Error_{phase} = \sum_{i=1}^K \begin{cases} [\pi^{-1} \Delta(\theta(T, T_i), \theta(C, T_i))]^2 & , T_i \in C \\ 1.0 & , T_i \notin C \end{cases} \left( \frac{K - i + 1}{K} \right) \quad (6.2)$$

Other intermediate measures created using separated phase objectives and various scaling schemes had been attempted with negligible, if not occasionally negative, performance changes when considering the set of targets as a whole.

Table 6.2: Experiment P1 Summary Table;  $K = 50$ 

| Target  | Agg.           | Power  |       | Phase |       |       |
|---|----------------|--------|-------|-------|-------|-------|
|   |                | Mean   | Best  | Mean  | Best  |       |
|    | Composition_01 | Mean   | 11.51 | 8.03  | 3.03  | 1.51  |
|   |                | StdDev | 1.03  | 1.25  | 0.42  | 0.30  |
|    | Composition_02 | Mean   | 5.90  | 0.00  | 2.34  | 0.00  |
|   |                | StdDev | 1.10  | 0.00  | 0.48  | 0.00  |
|    | Composition_03 | Mean   | 3.19  | 0.00  | 1.31  | 0.00  |
|   |                | StdDev | 1.22  | 0.00  | 0.60  | 0.00  |
|    | Composition_04 | Mean   | 6.28  | 0.43  | 2.65  | 0.06  |
|   |                | StdDev | 1.95  | 1.07  | 0.83  | 0.13  |
|    | Composition_05 | Mean   | 8.51  | 3.23  | 4.39  | 1.93  |
|   |                | StdDev | 1.72  | 1.72  | 0.82  | 0.72  |
|    | Composition_06 | Mean   | 12.04 | 6.90  | 5.05  | 2.90  |
|   |                | StdDev | 2.15  | 2.70  | 0.78  | 0.90  |
|   | Composition_07 | Mean   | 34.80 | 31.10 | 15.31 | 13.18 |
|   |                | StdDev | 1.62  | 3.21  | 0.96  | 1.58  |
|  | Composition_08 | Mean   | 36.91 | 33.03 | 16.29 | 13.85 |
|   |                | StdDev | 4.08  | 5.05  | 2.32  | 2.77  |
|  | Composition_09 | Mean   | 6.16  | 1.37  | 2.48  | 0.17  |
|   |                | StdDev | 1.52  | 0.89  | 0.80  | 0.19  |
|  | Composition_10 | Mean   | 26.08 | 19.50 | 14.36 | 11.43 |
|   |                | StdDev | 5.00  | 6.46  | 1.65  | 2.12  |
|  | Composition_11 | Mean   | 23.41 | 19.10 | 11.10 | 8.69  |
|   |                | StdDev | 2.69  | 2.84  | 1.41  | 1.56  |
|  | Composition_12 | Mean   | 33.02 | 28.33 | 15.77 | 13.19 |
|   |                | StdDev | 2.98  | 3.33  | 1.31  | 1.34  |

We appear to have obtained reasonable success for most of our targets with this approach (with difficulty remaining for Composition\_07 and Composition\_08). Inspecting the evolved examples in Figure 6.5, we see results largely similar to those found in experiment J3, with some minor improvements to Composition\_04, and variations of Composition\_07-08.

Targets Composition\_02-04 were found to be largely trivial with our previous J3 and K3 measures, and so they remain with P1. Results for Composition\_05-06 are of middling quality; not quite as ideal as those produced from the K1 measure, but still very recognizable, and further improved with lower choice of  $K$ . In Table 6.2,

Figure 6.5: Experiment P1 Compositional Summary Charts & Examples;  $K = 50$

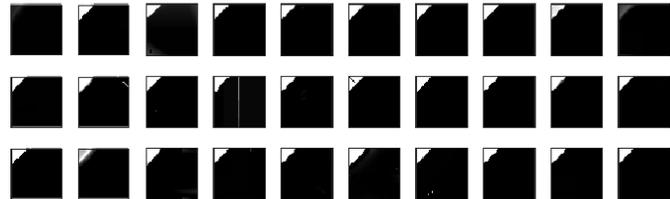
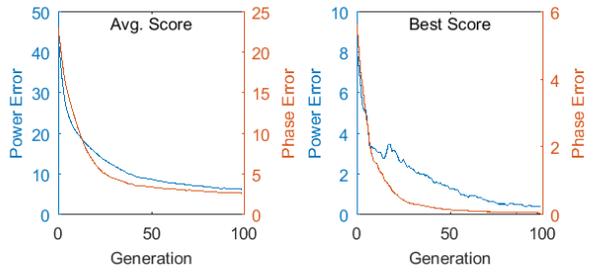
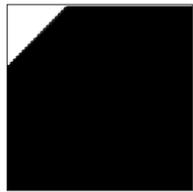
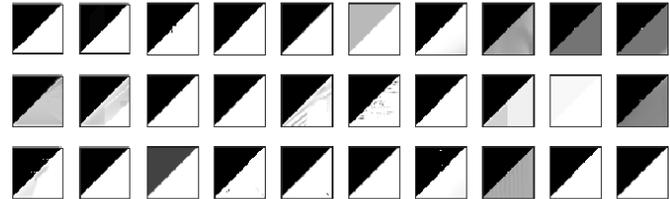
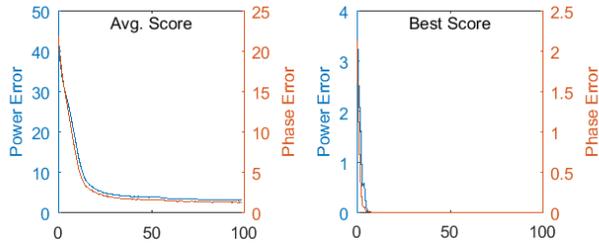
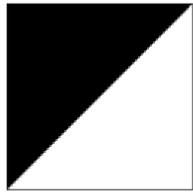
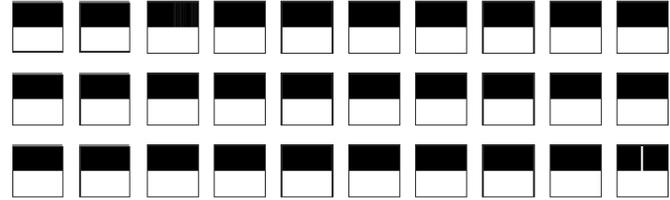
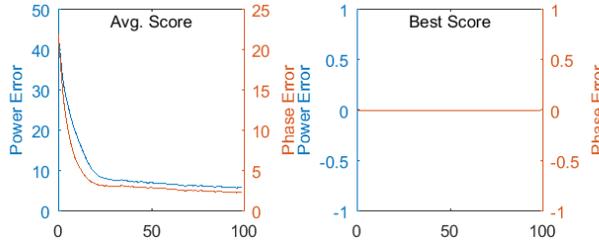
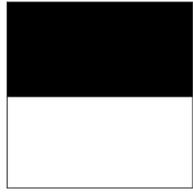
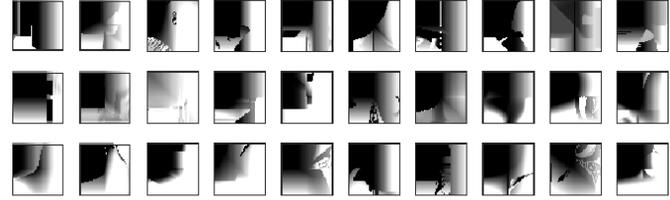
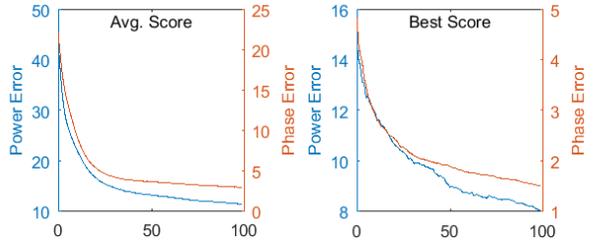
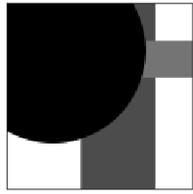


Figure 6.5: Experiment P1 Compositional Summary Charts & Examples;  $K = 50$

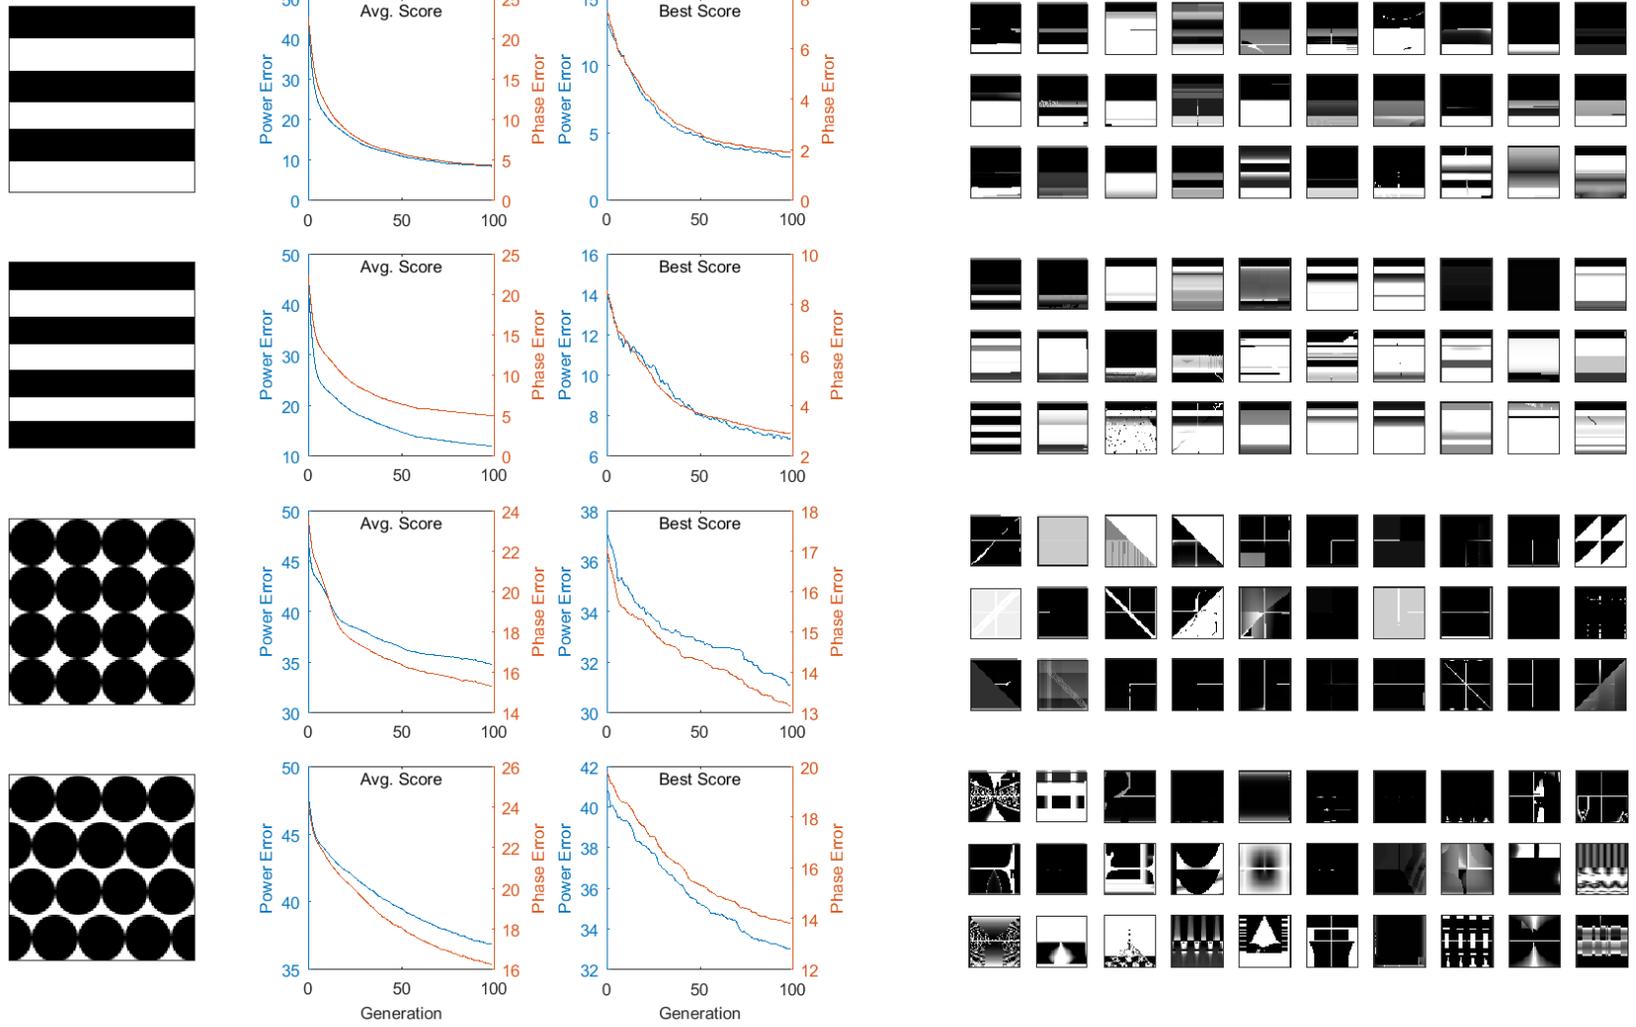
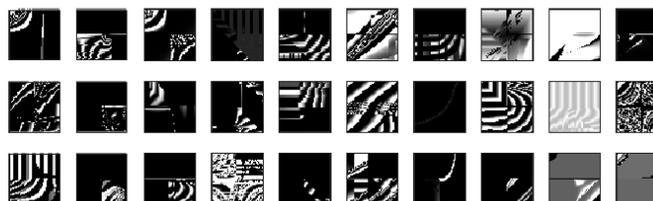
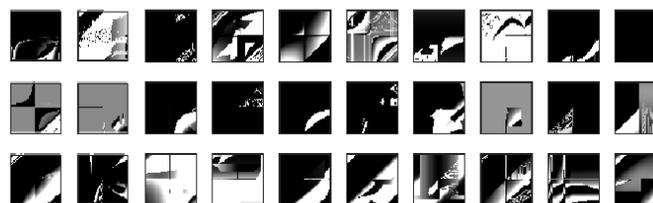
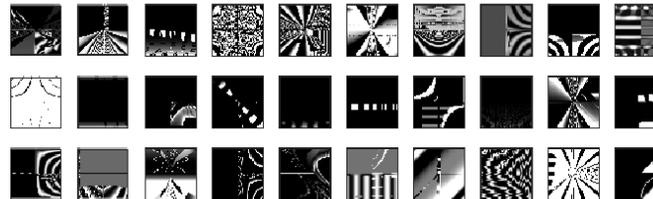
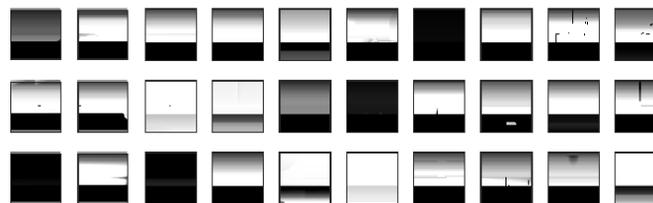
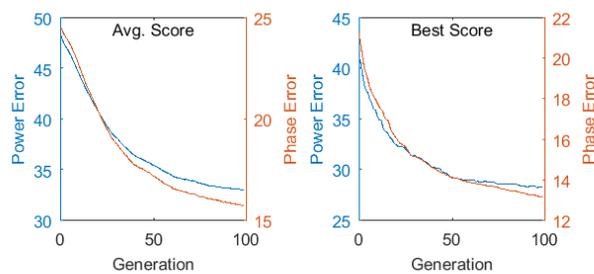
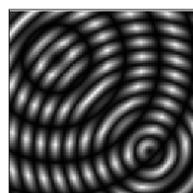
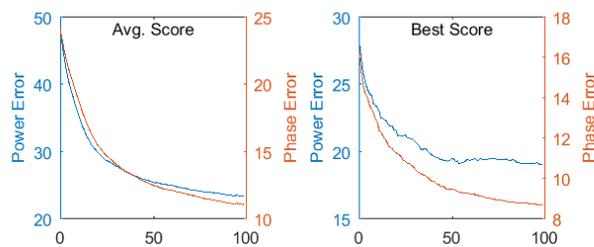
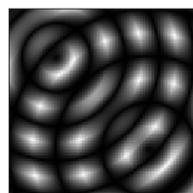
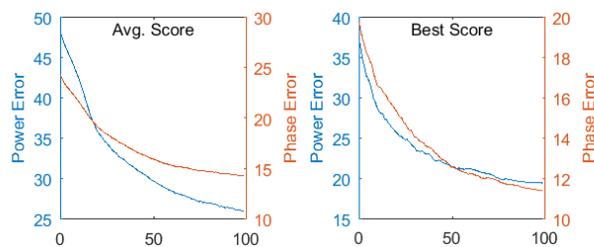
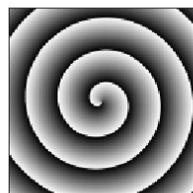
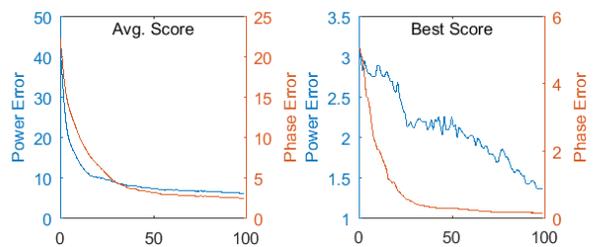


Figure 6.5: Experiment P1 Compositional Summary Charts & Examples;  $K = 50$ 

we see that the phase objectives score of Composition\_05-06 are comparatively lower than the power objective scores. With the difference in calculation methods for the two objectives, this is not unexpected, though the performance relative to some of the other outlier targets is a welcome change. Contrarily, the higher error in the power objective for Composition\_05-06 despite the improved phase scores perhaps corroborates the increased visual performance with lower  $K$  values. This is in line with what we found while re-evaluating J3 (Figure 6.4) on the new target set with  $K = 10$ , and we see similar results in the current experiment with  $K = 10$  or  $K = 25$ . The more artistic targets, Composition\_01 and Composition\_09, produce varied but reasonably recognizable recreations.

There is still an apparent difficulty when using any measure with targets Composition\_10-12 (spiral/Gabors), and especially Composition\_07-08 (circle repetitions). These targets show comparatively raised error values in Table 6.2 for both phase and error objectives. When previously revisiting experiment K3 with our compositional targets, we found excellent performance in the power objective scores, despite phase error remaining high. Visually, results showed similar contrast frequencies, but still often lacked compositional resemblance to the target. The current mediocrity of both P1 objective scores on these targets is still far from ideal, but we begin to suspect that our system is unable to further guide evolution of these targets' properties without compromising between power and phase. When we consider the shared spatial characteristics between these high-error targets, it seems that GP language limitations, and not fitness measures, could be unduly impeding candidates from producing circular traits. We consider it worth further investigation before any continued fitness refinements.

## 6.4 Survey Preparation

While larger adjustments to our GP language will be explored later in Section 8.1, observing the prolonged difficulty with targets Composition\_07,08,10,11,12 which feature circular characteristics has led to the thought that our basic GP language is inherently difficult for these traits. One minor adjustment which we were able to accommodate prior to beginning our user validation survey was to evolve and examine candidates through our J1 measure with a few additional coordinate variables in the GP language set.

We briefly define the additional variables in Section 4.2.1, Table 4.5, though further details will be outlined in 8.1. Our extension to the language set appends the two

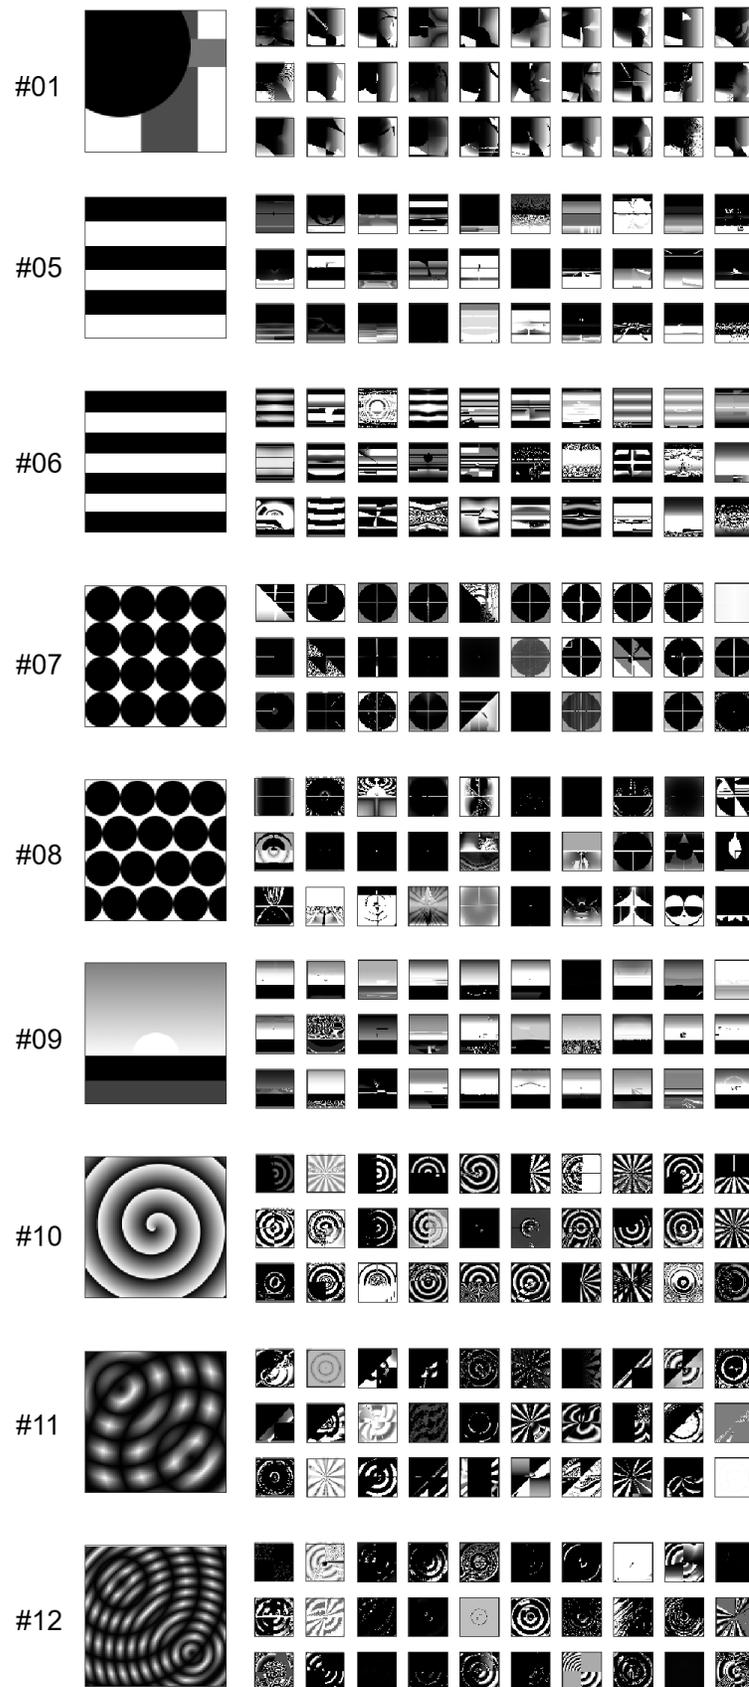
Table 6.3: Choice of  $K$  for Compositional Targets

| Target         | $K$ | Target         | $K$ | Target         | $K$ |
|----------------|-----|----------------|-----|----------------|-----|
| Composition_01 | 25  | Composition_05 | 25  | Composition_09 | 25  |
| Composition_02 | 10  | Composition_06 | 10  | Composition_10 | 50  |
| Composition_03 | 10  | Composition_07 | 50  | Composition_11 | 25  |
| Composition_04 | 10  | Composition_08 | 50  | Composition_12 | 25  |

polar coordinate variables,  $\rho$  and  $\phi$ , in addition to the existing normalized  $X$  and  $Y$  texture coordinates. We believe that this should reduce the difficulty for our system to present circular spatial aspects as needed.

As we have determined that selection of  $K$  should accommodate the specific target for improved results, we have used the previous P1 runs along with those produced with our expanded GP language to propose  $K$  values for our compositional target set. We detail the  $K$  values chosen per target in Table 6.3, and Figure 6.6 presents our candidates evolved using the new P1 measure, polar coordinate language set, and selected  $K$  values.

With these adjustments to our GP language, the spiral and Gabor targets (Composition\_10-12) show definite improvements regarding consistency and resemblance to their targets. However, difficulty is still seen in using Composition\_07 and Composition\_08 as targets. With the inclusion of polar coordinates, we begin to see some minor circular aspects, but our system with this measure is still having trouble pairing them with grid behaviour. We will see in Section 8.1 that some minor extensions to the GP language make these targets much more tractable. Section 8.1 will also further explore the performance effects of the polar coordinate language adjustments.

Figure 6.6: Polar Coordinate Language Summary Examples;  $K$  Per Target

# Chapter 7

## User Validation Survey

After having tested numerous refinements to our fitness measures, we found that some of the measures appeared to have enough utility for a general purpose compositional guide. However, as the compositional similarity of a pair of images can broach similar discussions as computational aesthetics, we found it of great importance to conduct a user survey and validate our findings with external observers. We need to see whether or not our fitness measures are in line with other human opinions. By conducting a user validation survey, it was hoped that we could remove some subjectivity regarding the success of the fitness measures, and determine if there was viewer confirmation of visual compositional similarity when using Fourier-guided fitness measures across various target images.

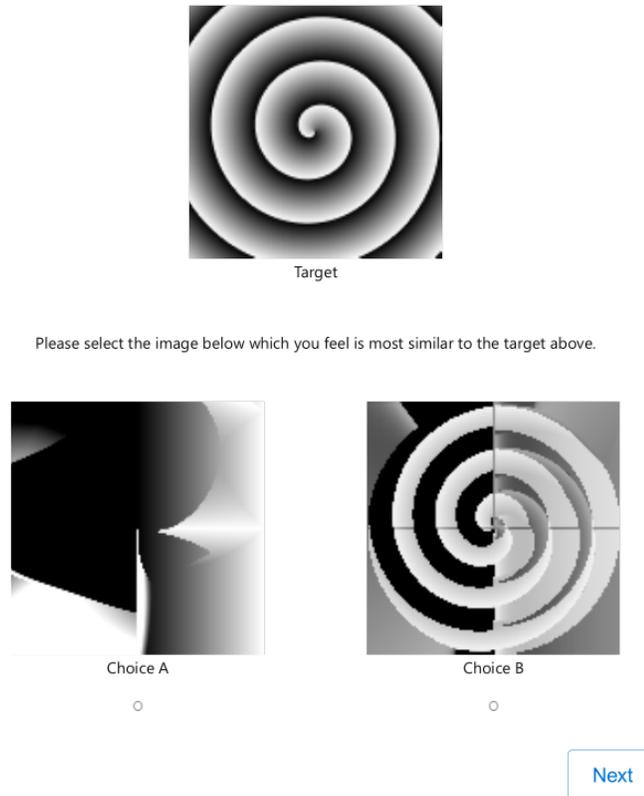
### 7.1 Survey Design

As we have declared a key objective for success in compositional guidance of our fitness measures to be the development of visually distinct groups, we must design a survey such that we can observe the rate of successful classification between some target, a matching generated texture, and a mismatching texture. We are hoping to verify whether or not survey participants believe that a texture generated from a GP run (using one of our developed fitness schemes) is more visually similar to the target image used for the run than some other texture generated for a different target.

Survey design resembles those used by Walsh *et al.*[65] and Salimi [66] in their similar use for user evaluations of fitness measures. The survey participants were presented with a target image (used to generate a “successful”, matching candidate). After observing the target, the user must determine which of two following images more closely matches this target. By recording the frequency at which user opinions

match the known production method (whether or not the selected image was produced with our fitness measures), we can find the degree of success for our measure across a number of sample pairings. Figure 7.1 shows one such sample pairing of images included in the online survey.

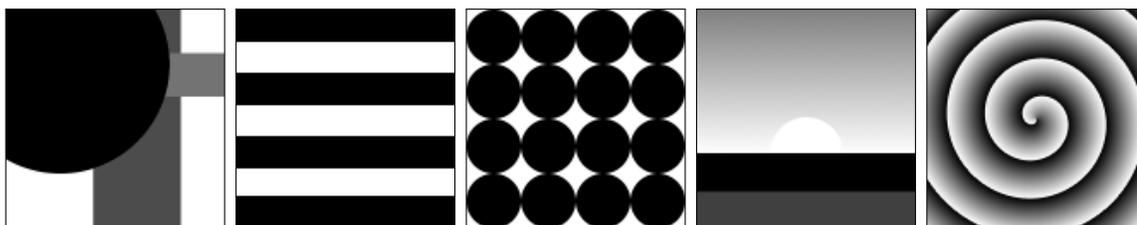
Figure 7.1: An example survey question. Choice B was evolved using the shown target image, where Choice A was evolved from a different target. Here, Choice B would be the correct match to our measure.



The target images used for the survey were obtained from a subset of the targets explored in Section 6.1. Where the entire set of targets include a number of purposefully similar compositional pairings, the targets for the user survey were to remain compositionally distinct. Figure 7.2 shows the collection of targets used in the survey. As the purpose of the survey was to validate spatial and compositional features, it was felt that including textures evolved using multiple colour channels could be distracting to the user, and thus monochrome targets and textures were used.

In designing the user survey, the goal of obtaining statistical results was balanced with maintaining user retention, as incomplete surveys would be discarded. It was decided that a limit of 20 questions would provide sufficient statistical power while not becoming a burden on the participating users. For each of the 5 target images,

Figure 7.2: Compositional targets for user survey



4 questions were displayed, each with a matching candidate produced by using the relevant target image, and a mismatching candidate which used one of the other targets while being produced. The pool of possible candidates, both matching and mismatching, were obtained from the *P1\_Polar* experiment set, generating 30 runs using K values of 10, 25, or 50 as appropriate for each target. The *P1\_Polar* experiment set used the P1 fitness variant with the inclusion of polar coordinate variables to the base GP language, which was found to produce reasonable results for all of the compositional testing target images. From this pool, a handful of obvious outliers were first filtered out by removing candidate images whose average intensity was in the top or bottom  $\sim 2\%$  of the intensity range (*i.e.* where the average pixel intensity was below 5, or above 250, on an 8-bit monochrome image). The choices of matching and mismatching images, placement of matching image on the questionnaire (A or B), and ordering of questions were all produced using random number generation.

A simple user survey was developed in the form of an online questionnaire, with hosting provided from the Brock University Department of Computer Science. The survey was open to all public participants for a duration of four weeks, with letters of invitation being explicitly sent to the senior students (BSc, MSc) of the Brock University Department of Computer Science and Department of Mathematics, as well as other acquaintances of the authors. All user survey materials have been cleared by the Brock University Research Ethics Board (File# 16-267 ROSS). A full printout of the survey presented to users, along with an additional analysis breakdown, has been included in Appendix C.

## 7.2 Results & Analysis

Upon closure of the online questionnaire, 36 complete response sets and 3 incomplete response sets were captured within the 4 week span. Incomplete responses were discarded due to the possible implication of withdrawn participation consent, and also data uniformity. An unexpectedly high success rate was found across the set of

questions, with the lowest scoring question having a successful match rate of  $\sim 60\%$ .

We employ both the sign test and Wilcoxon signed-rank test with a pre-determined confidence level  $\alpha = 0.05$ . The sign test has been used by Walsh *et al.*[65] and Salimi [66] in similar computer graphic user surveys. The signed-rank test can provide stronger statistical power, but consequently, is more difficult to satisfy, and assumes a symmetric distribution[67]. Upon collating the user survey data, the skew of the distribution was measured to check if the symmetry assumption for the signed-rank test was satisfied. The distribution skew was measured to be  $-2.225$ , which may impact the confidence of the signed-rank test.

However, one can also use the signed-rank test to determine the significance of a distribution skew about a median[68]. In comparing the skew about the median of  $\mu_0 = 0.5$  (random selection), a tailed signed-rank test provides significant reason to believe that the observed survey results are skewed favourably toward matching the fitness measure.

Concerning ourselves only with the improvement of successful image matches over random selection, we consider a right-tailed test when calculating sign and signed-rank. With the assistance of Matlab, and verification by [68], we found a significance both through the sign test with  $p \leq 0.000002$ , and through the signed-rank test with  $p \leq 0.00005$ , ensuring our required confidence of  $\alpha = 0.05$  was met.

A full listing of survey data, along with intermediate analysis breakdowns is displayed in Appendix C, Table C.1.

### 7.3 Conclusion

With 36 participants each answering 20 questions, we have found that there is a statistically significant match between our selected fitness measure and human compositional perception. While we make no claim at this time as to the feasibility of the fitness measure for all possible target images, we can see that when used against basic compositional targets, visually distinct classes of images can be produced. Our survey has validated that one of our proposed fitness measures sufficiently assesses the unique spatial and compositional properties of simple compositional targets, and is in-line with human perception with high statistical significance ( $p \leq 0.00005$ ).

# Chapter 8

## Artistic Exploration

We have reached the milestone of finding a measure capable of producing evolved images with visually similar spatial traits. Further exploration focuses on possible uses and performance of our measure for the production of evolutionary art. To this end, we will consider changes to language and colour scheme, and possible relations with aesthetics.

We will first consider enhancements and extensions to our GP language which may better reflect some of the more full-featured languages used for evolutionary art applications. We evaluate some possible multi-objective adaptations of our measures, expanding our capabilities from grayscale to coloured textures across multiple colour schemes. Based on some unanticipated observations in the upcoming explorations, we briefly consider a related measure shown in previously explored work which could lead to a possible new model of computational aesthetics.

### 8.1 Language & Representation

While a basic, trigonometric GP language was used for our initial exploration of fitness measures, the application of evolutionary art should require the inclusion of more creative language elements, such as noise or other geometric operators.

We had also begun to suspect, based on observations in our previous P1 experiment, that language restrictions were unduly limiting the possibilities of our system to produce visually similar results with our compositional targets. When reconsidering the initial genre target set (Section 5.1), we expect the effect of these restrictions to have been further exacerbated. Prior to our validation survey, we had considered including polar coordinates to our GP language and observed a notable improvement with certain target images containing radial aspects. We will step back briefly to

analyse the effects of these language changes, and check if further improvements can be seen when including additional, more-artistic GP language elements.

Moving forward, changes to the GP language will be cumulative; experiments with new language operators will also include all language elements explored up to that point. After the introduction of the polar coordinate variables, most experiments use a reduced run size of 9, due to the increased run times demanded from rendering the increasingly intensive candidate trees. Beyond the noted language changes and reduced run count, remaining parameters remain consistent to those listed in Section 4 (See Table 4.3).

An overview of the experiments found in this section with their summary tables and figures is outlined in Table 8.1.

Table 8.1: Overview of Section 8.1 Experimental Variations

| Label | Measure | K | Language    | Tables | Figures   | Notes  |
|-------|---------|---|-------------|--------|-----------|--|
| P1    | P1      | † | Polar       | 8.2    | 8.2       | Basic language with the inclusion of polar coordinate variables                                      |
|       | P1      | † | Circle      | 8.3    | 8.3       | All previously mentioned language elements, plus circle, grid, and offset operators                  |
|       | P1      | † | Noisy       | 8.4    | 8.4, 8.5  | All previously mentioned language elements, plus a variety of noise operators                        |
|       | P1      | † | Noisy, $-X$ | —      | 8.9, 8.10 | All previously mentioned language elements, except the $X$ coordinate variable is explicitly removed |

† Choice of K in these experiments varies per target. Selections are outlined in Table 6.3, and reiterated within the tables of this section.

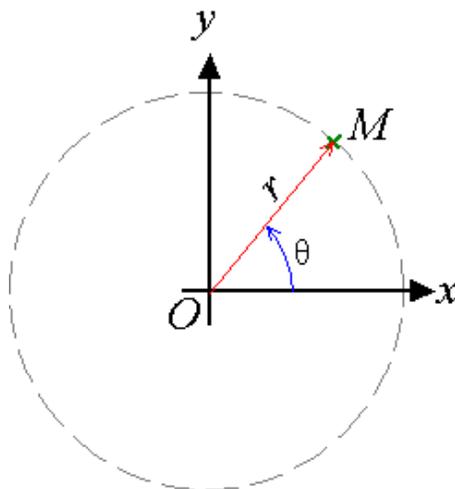
### 8.1.1 Polar Coordinates

The first adjustment to our GP language is motivated by the poor performance seen when using targets which displayed strong radial attributes. Analysis of results from the previous P1 experiment showed this commonality amongst poor performing targets, and so the inclusion of polar coordinates was considered as a means to reduce the difficulty of producing rough, radiant spatial characteristics.

In addition to the standard  $X$  and  $Y$  coordinate variables, we can compute from these the polar coordinates of a given position, as seen in Figure 8.1. By making these values available as coordinates, we may be more likely to see candidates with transitions across rotation, and radial fades or contrasts; behaviours that would be unlikely to appear without stumbling across trigonometric operators and coincidentally reproducing the polar coordinate transforms. As these are precisely the characteristics that we are hoping for in the high-error targets, we were hopeful that this language change would resolve the remaining evolutionary difficulties.

Figure 8.1: Polar Coordinate System [69]

The point  $M$  can be represented by using the pair of distance/radius ( $r$ ) and radial angle ( $\theta$ ) about the origin. Alternative naming conventions use  $\rho$  and  $\phi$  respectively.

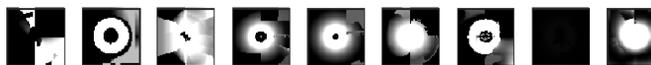
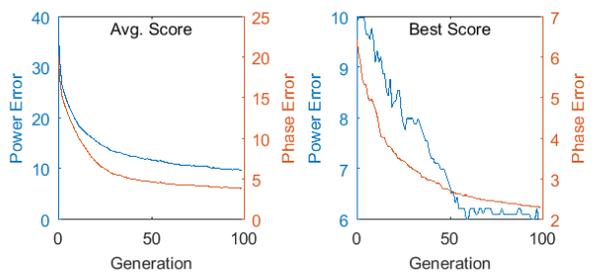
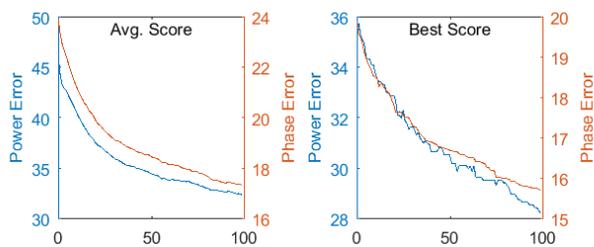
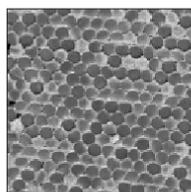
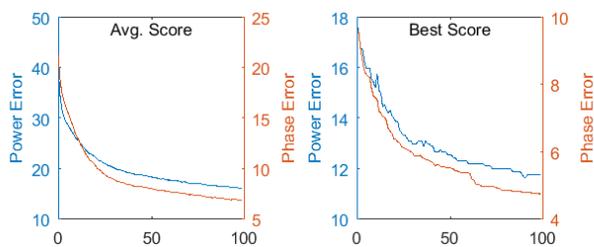
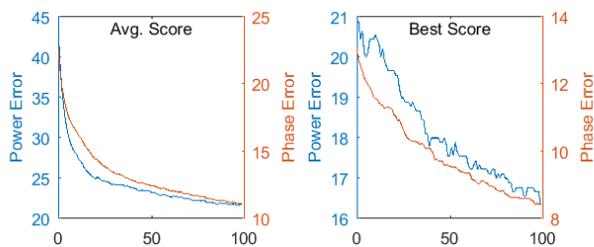
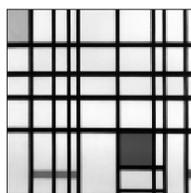


Examples of candidates evolved using the polar coordinate set have been seen previously with Figure 6.6 while preparing the user survey, and visual inspection of the difficult targets suggested a strong partial success. We observed that the later Composition\_10-12 showed notable improvements, where some issues still remain with Composition\_07-08. When we return to the initial genre target set, Figure 8.2 shows a few small but notable changes. With the flower target, we see a central radiant shape

Table 8.2: Polar Language Summary Table

|   | Target         | $K$ | Agg.           | Power         |               | Phase         |               |
|---|----------------|-----|----------------|---------------|---------------|---------------|---------------|
|   |                |     |                | Mean          | Best          | Mean          | Best          |
|    | Composition_01 | 25  | Mean<br>StdDev | 2.34<br>0.68  | 0.40<br>0.81  | 0.79<br>0.18  | 0.08<br>0.05  |
|    | Composition_02 | 10  | Mean<br>StdDev | 6.83<br>1.05  | 0.00<br>0.00  | 2.65<br>0.48  | 0.00<br>0.00  |
|    | Composition_03 | 10  | Mean<br>StdDev | 3.27<br>1.19  | 0.00<br>0.00  | 1.35<br>0.55  | 0.00<br>0.00  |
|    | Composition_04 | 10  | Mean<br>StdDev | 6.32<br>1.52  | 1.20<br>1.40  | 2.12<br>0.63  | 0.08<br>0.12  |
|    | Composition_05 | 25  | Mean<br>StdDev | 8.45<br>1.56  | 5.93<br>1.86  | 3.17<br>0.75  | 1.61<br>0.76  |
|    | Composition_06 | 10  | Mean<br>StdDev | 12.45<br>2.92 | 6.80<br>3.21  | 5.51<br>1.13  | 3.15<br>0.99  |
|   | Composition_07 | 50  | Mean<br>StdDev | 34.36<br>3.26 | 30.47<br>4.44 | 14.31<br>1.16 | 12.09<br>1.46 |
|  | Composition_08 | 50  | Mean<br>StdDev | 37.15<br>3.07 | 32.67<br>4.74 | 16.35<br>1.74 | 13.81<br>2.51 |
|  | Composition_09 | 25  | Mean<br>StdDev | 2.79<br>0.78  | 0.13<br>0.51  | 1.17<br>0.36  | 0.08<br>0.13  |
|  | Composition_10 | 50  | Mean<br>StdDev | 15.79<br>4.82 | 9.50<br>4.66  | 8.51<br>2.89  | 5.51<br>3.04  |
|  | Composition_11 | 25  | Mean<br>StdDev | 11.53<br>2.51 | 7.40<br>2.74  | 5.51<br>1.08  | 3.44<br>1.15  |
|  | Composition_12 | 25  | Mean<br>StdDev | 14.59<br>1.85 | 11.20<br>1.71 | 6.33<br>0.91  | 4.65<br>0.76  |
|  | Mondrian       | 50  | Mean<br>StdDev | 21.80<br>2.32 | 16.33<br>2.83 | 11.18<br>1.26 | 8.42<br>1.13  |
|  | Van Gogh       | 50  | Mean<br>StdDev | 16.24<br>1.35 | 11.78<br>0.83 | 6.93<br>0.65  | 4.76<br>0.45  |
|  | Cable Ends     | 50  | Mean<br>StdDev | 32.42<br>1.93 | 28.22<br>2.82 | 17.35<br>0.89 | 15.72<br>1.23 |
|  | Flower         | 50  | Mean<br>StdDev | 9.87<br>0.80  | 6.00<br>0.71  | 3.91<br>0.93  | 2.31<br>1.01  |

Figure 8.2: Polar Coordinate Language Genre Targets Summary Charts & Examples



with similar proportion to the flower petals in the target. With the Van Gogh target, while still rough, we see a rather consistent return to a top-body portrait figure. A few examples appear to be attempting to capture traits resembling a head.

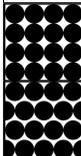
A summary of all compositional targets using  $K = 50$  can be found in Appendix B Table B.1, but as we have found that each target shows improved performance when matched to a specific  $K$  value, we will use these specific target-dependent  $K$  values for Table 8.2 and also for when we compare the upcoming language changes. While we can occasionally see some strong visual indicators that performance has improved, we will also perform various 2-sample t-tests on the population means between language adjustments to see if there have been any statistically significant changes beyond our perception. We will note, of course, the unfortunate requirement to reduce run size and the subsequent reduction of statistical power.

While we had observed improvements to the flower and Van Gogh targets from the initial genre set, we also see significant reductions to their phase error scores, with  $p < 10^{-19}$ . Power error scores did show a statistically significant change, though not to the same extent, with the flower target giving a slight improvement ( $p = 0.02$ ), and the Van Gogh actually having a minor deterioration ( $p = 0.03$ ). With the compositional target set, we had noted strong improvements to Composition\_10-12, and this is reflected statistically as improvements to both fitness objectives ( $p < 0.00001$ ). While it is unfortunate that the circle grid targets (Composition\_07-08) did not substantially improve, we found that Composition\_07 actually showed significant improvement with regard to phase error ( $p < 0.0007$ ). However, there was a substantial amount of error to overcome, and the improvement — while significant — is still wanting. Composition\_01, which had performed quite well previously, found a slight but significant improvement. Contrasting our improvements, we can observe a statistically significant degradation in performance with Composition\_05 (stripes). While both stripe targets performed better without the polar coordinates, Composition\_05 gives a notable loss with  $p < 0.002$ . This was not altogether unexpected, as the addition of operators intrinsically tied to rotations is poorly suited for the straight stripe targets, and a slight negative language bias should be considered.

### 8.1.2 Circle, Grid, & Offset

The inclusion of polar coordinate variables has improved results for certain compositional targets and some of our initial, genre image targets. Despite this, we see that some of the seemingly simple, circular grid targets are still having difficulty re-

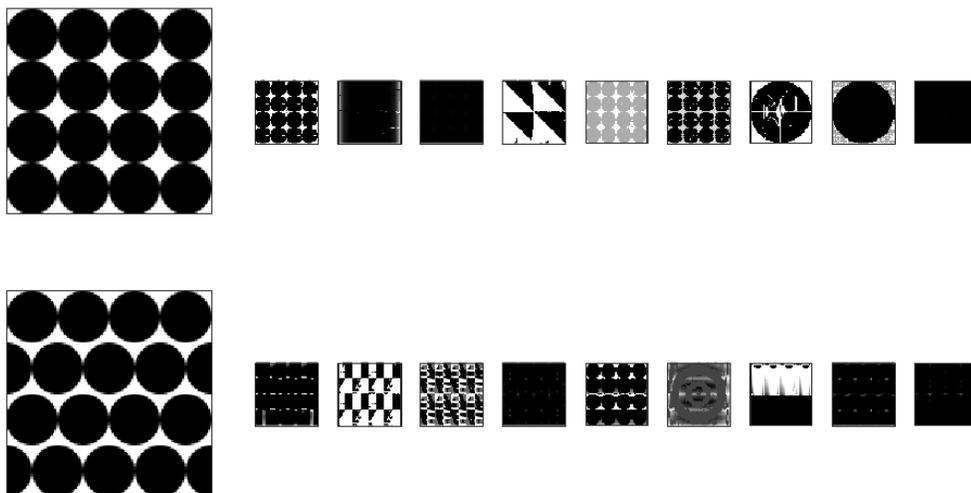
Table 8.3: Circle, Grid, Offset Language Summary Table

| Target  | $K$ | Agg.   | Power |       | Phase |      |
|---|-----|--------|-------|-------|-------|------|
|   |     |        | Mean  | Best  | Mean  | Best |
|  | 50  | Mean   | 22.51 | 18.00 | 9.44  | 7.28 |
|   |     | StdDev | 13.52 | 16.13 | 5.94  | 6.55 |
|  | 50  | Mean   | 23.91 | 17.00 | 10.59 | 7.41 |
|   |     | StdDev | 6.56  | 8.46  | 2.92  | 3.54 |

producing their basic spatial aspects. Before moving forward with the multitude of noise operators, we will first give one last attempt at recreating the difficult targets of Composition\_07 and Composition\_08.

To improve the capabilities of our system and give the best possible chance of success, we will add a set of GP language operators well-suited for these particular targets. With inspiration from the Gentropy system by Weins [44], we will include 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 provides a simplified way for the candidate programs to show hard transitions about a radius, and the *tile* operator provides an easy way to create arbitrary  $n \times m$  tilings about a grid. These two operators should allow Composition\_07 to see much more success, where the additional *shift* operator (being able to offset the texel window for its children) should greatly assist in producing the alternating offset aspects in Composition\_08.

Figure 8.3: Circle, Grid, Offset Language Summary Examples



We try these language changes for the two target ( $K = 50$ ), and are pleased with the results. Figure 8.3 shows a much-improved set of evolved candidate tex-

tures over our previous experiments. We see result categories split into approximate thirds between featureless outliers, rough variations, and high-similarity reproductions. In addition to the strong visual improvements, when we compare Table 8.3 to the previous Table 8.2, we see the error for these two targets decreases by  $\sim 40\%$  in both objectives. 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 language adequacy. As these targets were one of the most difficult categories used in the user survey, it is unfortunate that these results could not be obtained in time for that purpose, though we see that survey results were still quite satisfactory.

### 8.1.3 Noise Generation

As we have now reached satisfactory results with our compositional targets, we will return our focus to the initial, more-artistic genre targets. To help facilitate our goal of producing evolutionary art, we will finally end our current language exploration with the addition of numerous noise generation operators. Details regarding noise operators can be found in Section 4.2.1. Here we begin to see significant increases to run times.

Composition\_01 and Composition\_11 (in Figure 8.4) showed slight improvement with the addition of the noise operators ( $p < 0.02$ ), as did – surprisingly – Composition\_02 ( $p < 0.0002$ ). We suspect that the increased performance with Composition\_02 may be due to noise operators having higher likelihood to be introduced through mutation in the more complex candidate trees. With Composition\_02 requiring only a simple expression to produce an ideal candidate, and with noise being likely to express power in many non-target coefficient positions, it may be effectively adding additional selective pressure against the complex candidates. Composition\_07-08 showed marked improvements over their polar counterparts, but no substantial changes from the previous evaluations with *circle* and *tile* operators included (noting that the previous *circle*, *tile*, and *shift* operators are included in the noisy language set). The other compositional targets showed either no significant improvement, or a significant degradation, when including noise operators in their language set.

When we return to the results produced for the genre targets (Figure 8.5), we see that the candidate images for the flower target have substantially improved. The

Figure 8.4: Noisy Language Summary Examples; Compositional Targets;  $K = 50$   
Mixed results were reported, with statistical analysis suggesting (from top to bottom)  
improvement ( $p < 0.02$ ), degradation ( $p < 0.05$ ), improvement ( $p < 0.02$ ), and no  
significant changes.

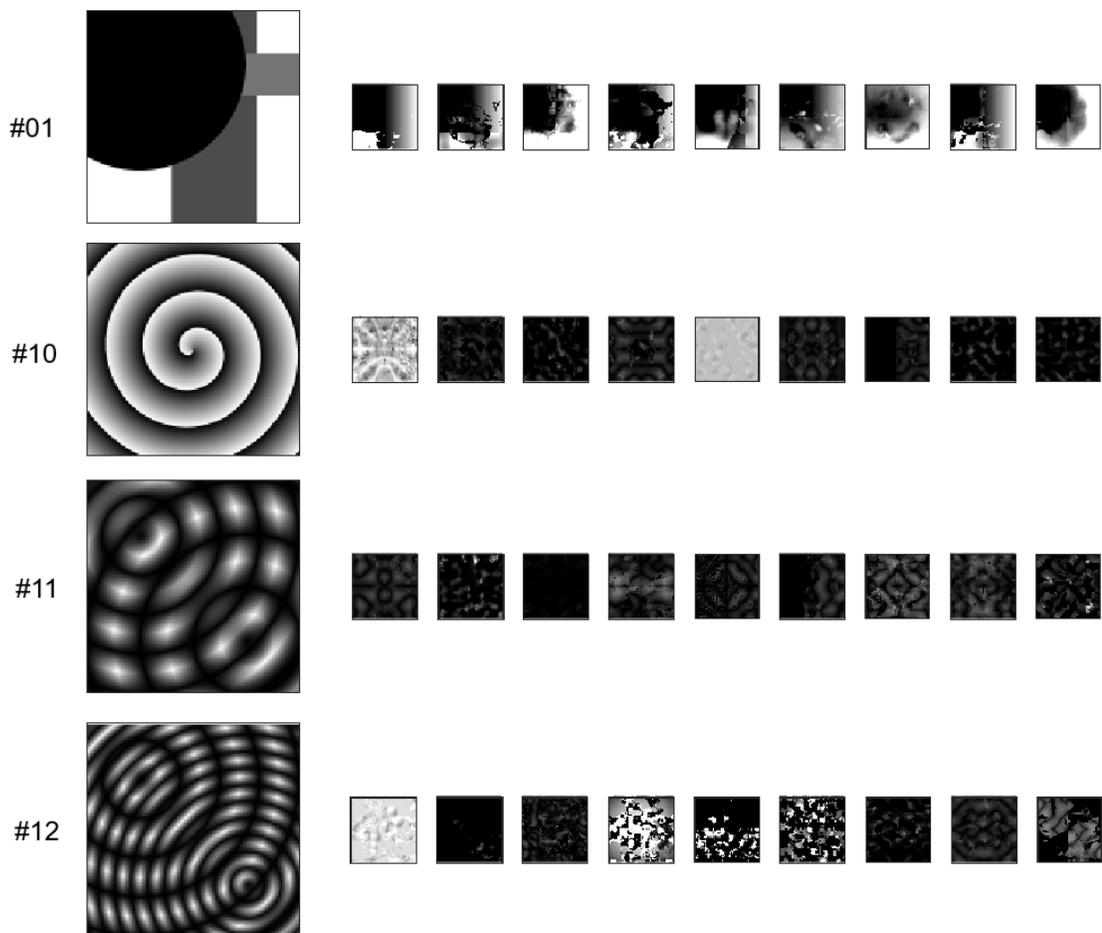


Figure 8.5: Noisy Language Summary Charts & Examples; Genre Targets;  $K = 50$

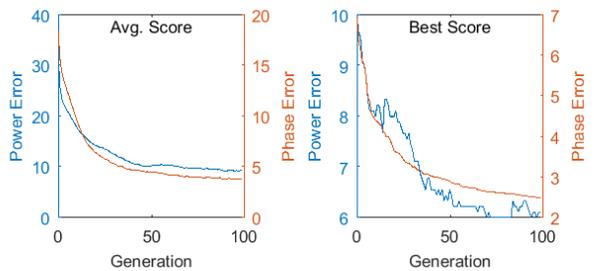
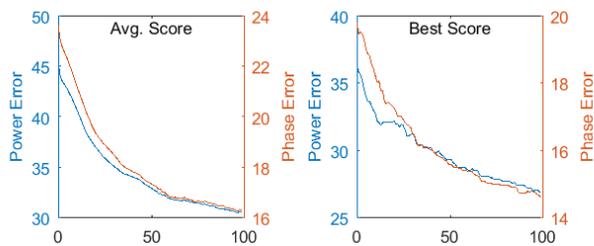
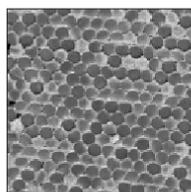
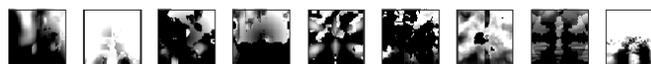
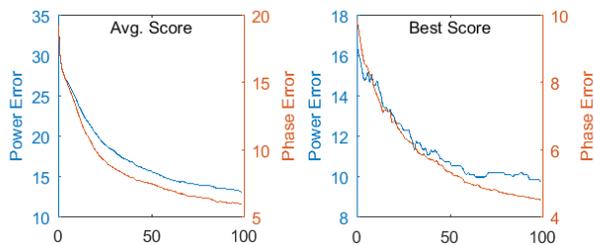
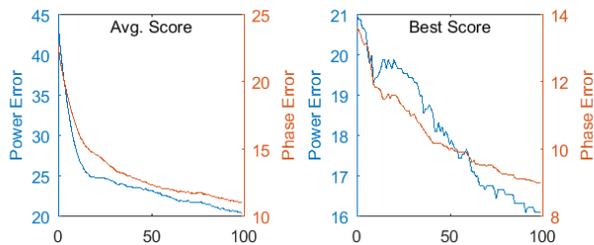
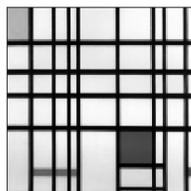


Table 8.4: Noisy Language Summary Table

|   | Target     | Agg.   | Power |       | Phase |       |
|---|------------|--------|-------|-------|-------|-------|
|   |            |        | Mean  | Best  | Mean  | Best  |
|  | Mondrian   | Mean   | 20.44 | 16.11 | 11.05 | 9.00  |
|   |            | StdDev | 1.37  | 2.03  | 1.07  | 1.71  |
|  | Van Gogh   | Mean   | 13.13 | 9.78  | 5.99  | 4.52  |
|   |            | StdDev | 2.95  | 3.42  | 0.73  | 0.82  |
|  | Cable Ends | Mean   | 30.63 | 26.89 | 16.34 | 14.60 |
|   |            | StdDev | 2.52  | 3.37  | 0.84  | 1.09  |
|  | Flower     | Mean   | 9.34  | 6.11  | 3.81  | 2.50  |
|   |            | StdDev | 1.69  | 1.36  | 0.40  | 0.69  |

Mondrian and Cable Ends targets at first glance have few visual differences from their noise-less renditions, though a closer investigation shows some finer details and frequencies can be found in what seemed to be plain, empty regions. The Van Gogh appears to have captured some of the traits it saw previously but has them now seemingly buried under an emergent busy and noisy layer.

Numerically, the introduction of the noise operators appears to be an improvement for these targets, with minor but consistent reductions in both phase and power errors. However, these changes are not largely reinforced through any statistical significance. After the substantial and significant changes seen with the flower target when moving from the base language to the polar coordinate extension, we had suspected that the current visual changes would also warrant similar numerical significance, yet it is not the case. We then begin to suspect that the current solutions have been able to trade sharpness for positional details. Solutions using only the polar coordinate language set tend to show stronger contrasts, like the flower and background, where the current noise-based solutions have weaker contrasts but the benefit of expressing individual petals. Conceivably, solutions could be biased toward placing more power at difference coefficient positions, yet still maintain similar power and phase error scores. These language biases may work favourably when pairing the natural flower image with noise operators (noted for conveying more natural aesthetics), or perhaps unfavourably, as seen with the current Van Gogh renditions.

Still, while a little different than initially expected, we are beginning to see traits and candidates that have come to expect with evolutionary art. Figure 8.6 highlights some of the finer details in a pair of larger renderings from current experiment for the flower target.

Figure 8.6: Produced Image Highlights; Noisy Language; Flower



#### 8.1.4 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 problems encountered while lacking the polar coordinate variables, there were few noted changes to performance. We do see a few small artifacts of this change through some of the traits commonly encountered in these resultant images.

With our artistic genre targets, the flower image appears to have returned to some of the stronger contrast results seen when using the polar coordinate set, and the Van Gogh may be slightly more abstract and noisy (Figure 8.9). No statistically significant changes to performance can be found. 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. Composition\_05-06 (the horizontal stripes) showed a slight improvement with the absence of the  $X$  coordinate variable, but not enough to provide statistical significance with a run count of 9.

When we inspect the evolved textures a little more closely, there appears to be two main ways that our system and its textures have adapted to the missing coordinate variable. Some candidates were able to glean sufficient positional information from the remaining coordinate variables:  $Y$ ,  $\rho$ , and  $\phi$ . We see this in some of the more circular

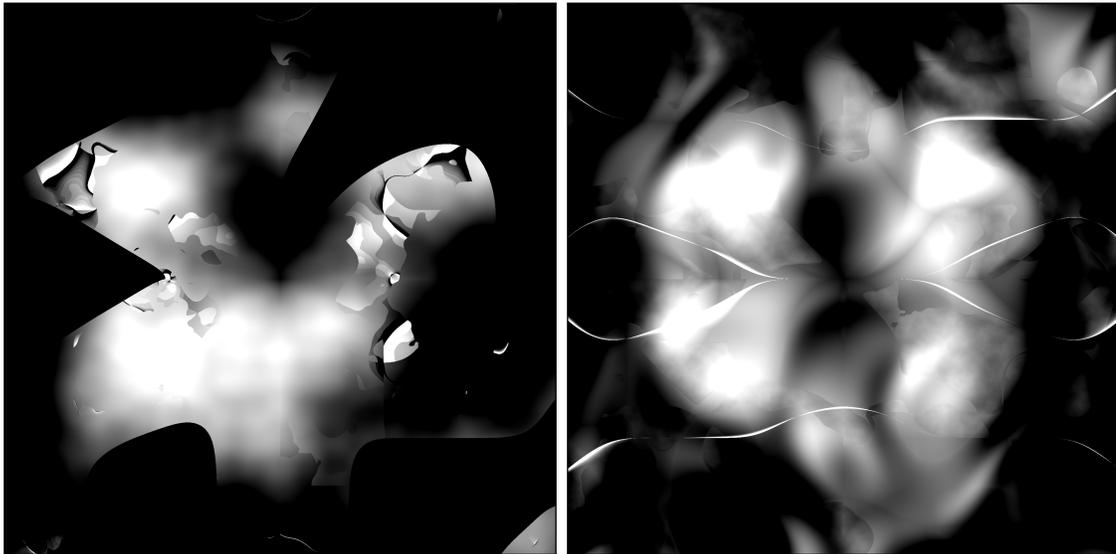
Figure 8.7: Produced Image Highlights; Noisy Language, No  $X$ ; Flower

Figure 8.8: 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.

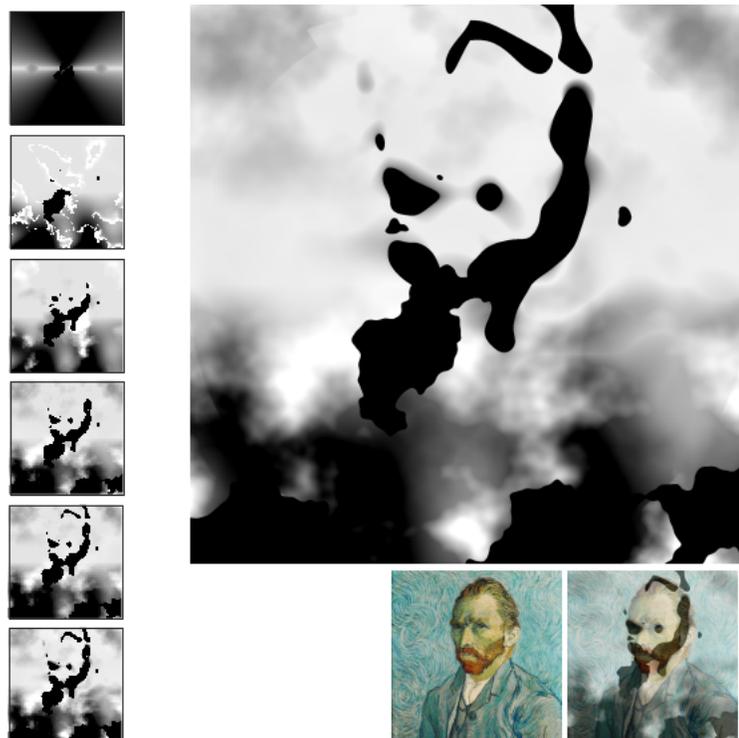


Figure 8.9: Noisy Language, No  $X$  Summary Charts & Examples; Genre Targets;  $K = 50$

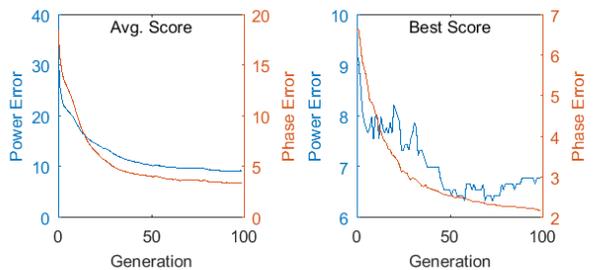
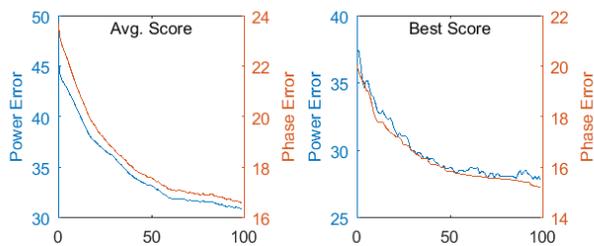
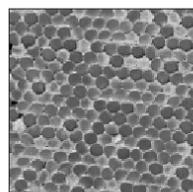
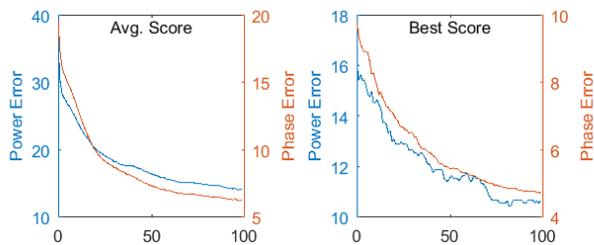
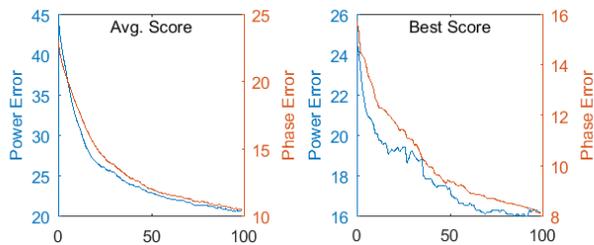
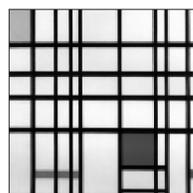
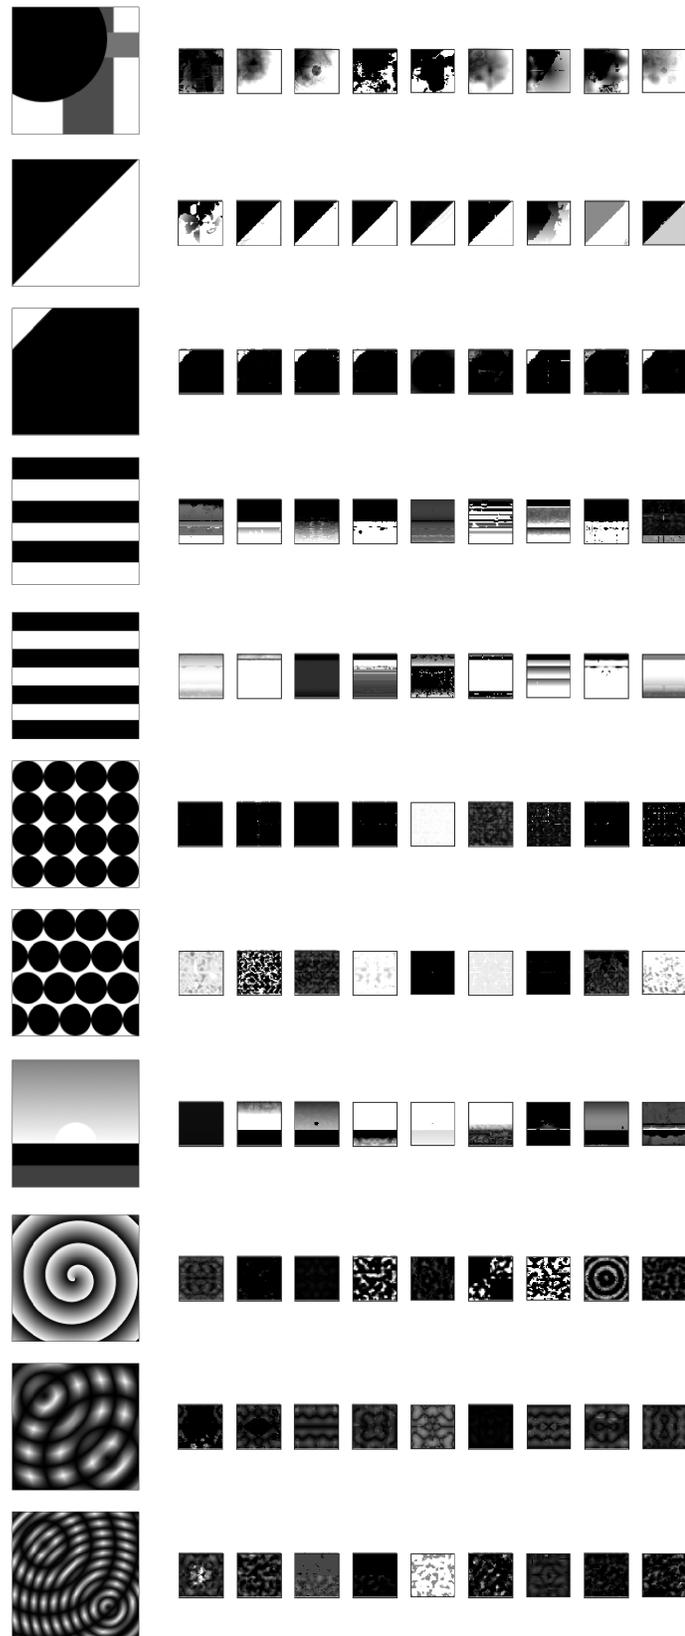


Figure 8.10: Noisy Language, No  $X$  Summary Examples



results from the flower target in Figure 8.9, and Composition\_01, Composition\_04, and Composition\_09 in Figure 8.10.

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 8.7, and the Van Gogh target in Figure 8.9. A particularly notable example of this is explored in Figure 8.8. Interestingly, we see this behaviour in the previously difficult circular targets; despite depending on the polar coordinates, it would seem that there is still a requirement for the basic  $X$  coordinate as well.

## 8.2 Colour

While we have seen some success in the production of similar grayscale candidate images, evolving useful textures often requires consideration of colour similarity. We will explore how capable our proposed fitness measure can be with regards to evolving images with similar colour distribution.

When we consider the concept of frequency analysis for colour images, we often run into problems trying to capture intensity changes running between colour channels [33]. Changes to average luminance cannot capture all colour changes, and spatial features can be composed across multiple colour channels. When considering an alternate colour model like HSL, we may run into issues with wrapping artifacts in the hue channel.

The two main differentiating techniques we explore will be in using a separate objective to reduce colour differences, and having colour channels receive their own individual objectives for reduction of power spectra error in their channel.

An overview of the experiments explored in this subsection can be found in Table 8.5.

Table 8.5: Overview of Section 8.2 Experimental Variations

| Label | Measure | K  | Language    | Tables | Figures | Notes  |
|-------|---------|----|-------------|--------|---------|--|
| C1    | P1      | 50 | Base        | —      | 8.12    | Measure applied to RGB channels separately                                       |
| C2    | P1      | 50 | Base        | —      | 8.13    | Measure applied to RGB and luminance (Y) channels separately                     |
|       | P1      | 50 | Noisy       | —      | 8.14    |  |
|       | P1      | 50 | Noisy, $-X$ | —      | 8.15    |  |
| C3    | P1      | 50 | Base        | —      | 8.17    | Measure applied to HSL channels separately                                       |
| C4    | P1      | 50 | Base        | —      | 8.18    | Measure applied to luminance channel, CHISTQ comparison used for colour matching |
|       | P1      | 50 | Noisy       | —      | 8.19    |  |
|       | P1      | 50 | Base        | —      | 8.20    | CHISTQ used to guide levels of graytones in monochrome textures                  |

### 8.2.1 Colour Targets

In performing our experiments with colour, we will use the original genre targets, but will supplement them with additional compositional images which have been adapted to display colour properties. We had seen previously that the genre targets were difficult to reproduce even when limited to grayscale, and so we suspect that we may still see difficulties in reproducing shape characteristics with these targets. However, they provide a suitable range of expected colour distributions for use with evolutionary art.

We will borrow *Composition\_01*, as it originally was produced with colour experimentation in mind. The additional compositional targets have been created by taking variations of the previously seen *Composition\_02-12* and using their intensity map to determine values for one of the red, green, or blue (RGB) colour channels of a new target. Figure 8.11 displays the new colour targets along with a decomposition into their RGB colour channels. We can see that individually, the target colour channels are familiar from our previous grayscale experiments, and had shown a decent measure of success in reproducing spatial aspects. The last target applies some additional offsets to the image in an attempt to slightly increase evolutionary difficulty.

Between the original genre targets, and these newly created colour, compositional targets, we expect there to be a fair range of difficulty for our colour experiments across targets.

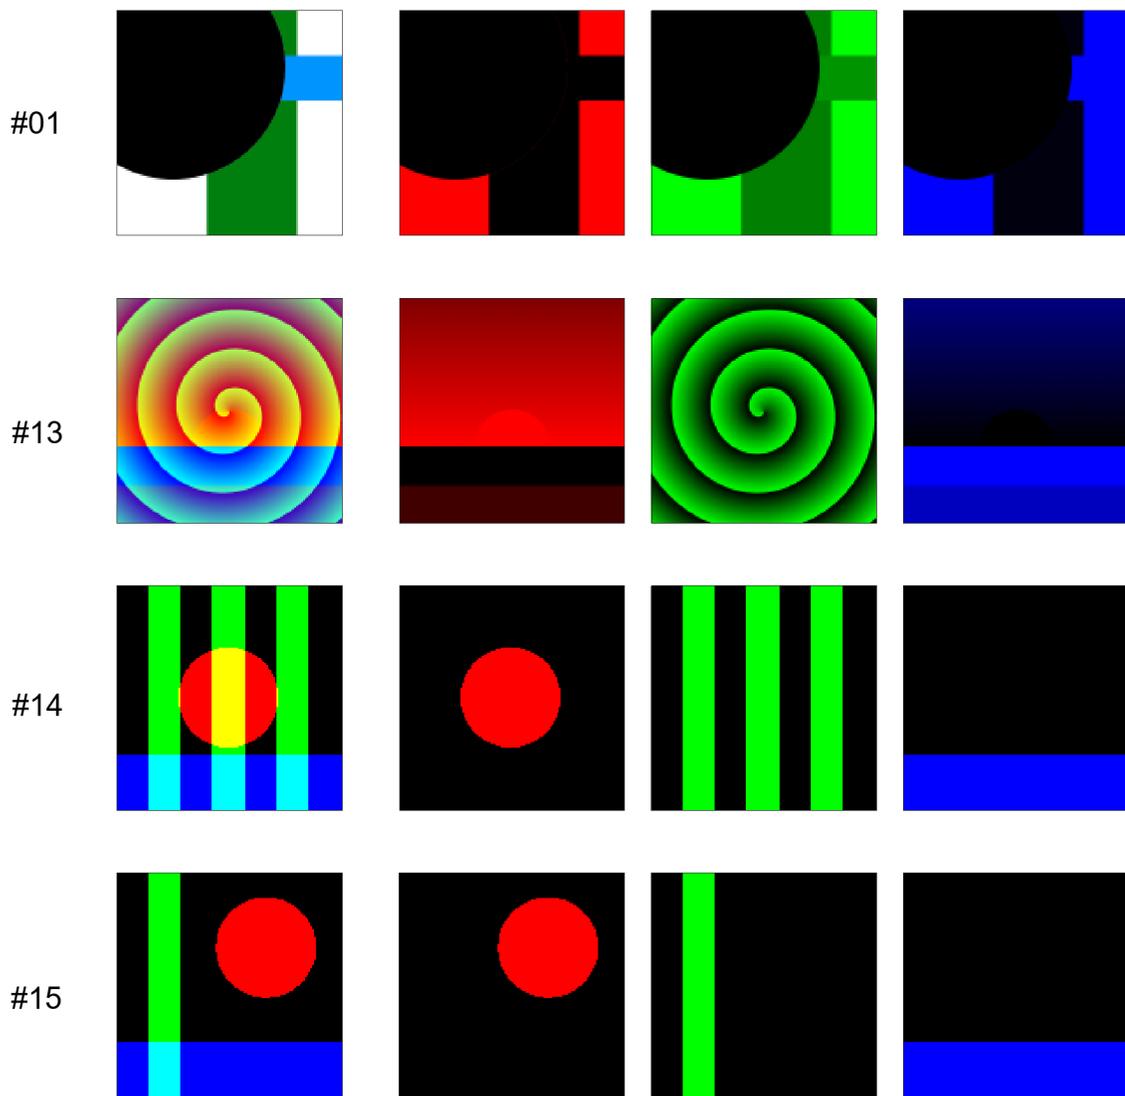
### 8.2.2 Objectives Per Colour Channel

We will continue our exploration using the system described previously. As we had seen language sufficiency issues with targets showing radial spatial aspects, we will include the polar coordinate variables in addition to the base GP language. We maintain the selection of  $K = 50$  as it produced suitable, albeit not necessarily optimal, results for the compositional targets, and was the preferred selection for the genre targets. To produce colour images, we will evolve three GP trees per individual, corresponding to the RGB colour channels. With the increased tree count, and proportional increase in rendering complexity, we will maintain our reduced run count of 9 runs per target.

As our GP system is configured to construct candidate images across trees for separate RGB colour channels, it is expected that these new targets will display improved visual results over the more complex genre targets. We have seen suitable evolutions using our system for the grayscale renderings, however the expansion to

Figure 8.11: Colour Target Set

The coloured, compositional targets are seen in the leftmost column, followed by their isolated red, green, and blue colour channels.



colour textures will still pose some difficulties for these otherwise basic targets. We cannot expect the same ease as evolving grayscale targets, as we will see an increase in potentially conflicting fitness objectives. We will maintain the same finite population counts and generation limits as the grayscale experiments, which effectively gives us a harder problem with the same amount of resources. We will later briefly discuss results using an extended generation count.

### **Colour Channels: RGB**

Our first expansion to colour textures will be to apply our best found all-purpose measure (P1) to each colour channel individually. Where a grayscale texture had two objectives (power and phase), our 3-channel (RGB) colour image will use  $2 \times 3 = 6$  objectives. Each channel will be evaluated similarly to a separate grayscale texture.

We expect that any differences between target and candidate images in a given colour channel will manifest as different, unexpected colour artifacts at that position. In the general case, colour dissimilarities between target and candidate will suggest an amount of error. However, grayscale evolutions that were considered successful have often seen similar, but still notably different, spatial composition; exacerbated further with creative and novel reproductions. Therein, this may be a potential limitation of these channel-independent approaches.

The results in Figure 8.12 show some promise with our new coloured compositional targets. With *Composition\_01*, we see that the large black section in the top-left corner – a feature common to all of its colour channels – is consistently reproduced. Although the colour channels for this target appear rather similar with only minor variation, it would seem that the green channel, having a more uniform intensity level, was favoured over the others. We can see spots of green appear in most of the candidates as per the target image, though reds and blues ended up being less constrained, as indicated by the frequent occurrence of magenta in the evolved candidates.

The remaining compositional images performed to our positive expectations. The performance of each colour channel appears to roughly coincide with the amount of difficulty seen when evolved as separate grayscale textures. One interesting observation is in the sacrifice of power error when evolving for *Composition\_15*. Despite having an identical blue channel target as its predecessor *Composition\_14*, we see elevated blue channel power error which seems to be exchanged for a reduction of green channel power error. This might suggest that the green channel is more volatile; having a wider range of results amongst the candidates to more strongly influence the

Figure 8.12: Experiment C1 Summary Charts & Examples

Power and phase errors for each of the red, green, and blue colour channels are plotted with their respective colour. Power error is denoted with solid lines, where phase error uses a dashed line.

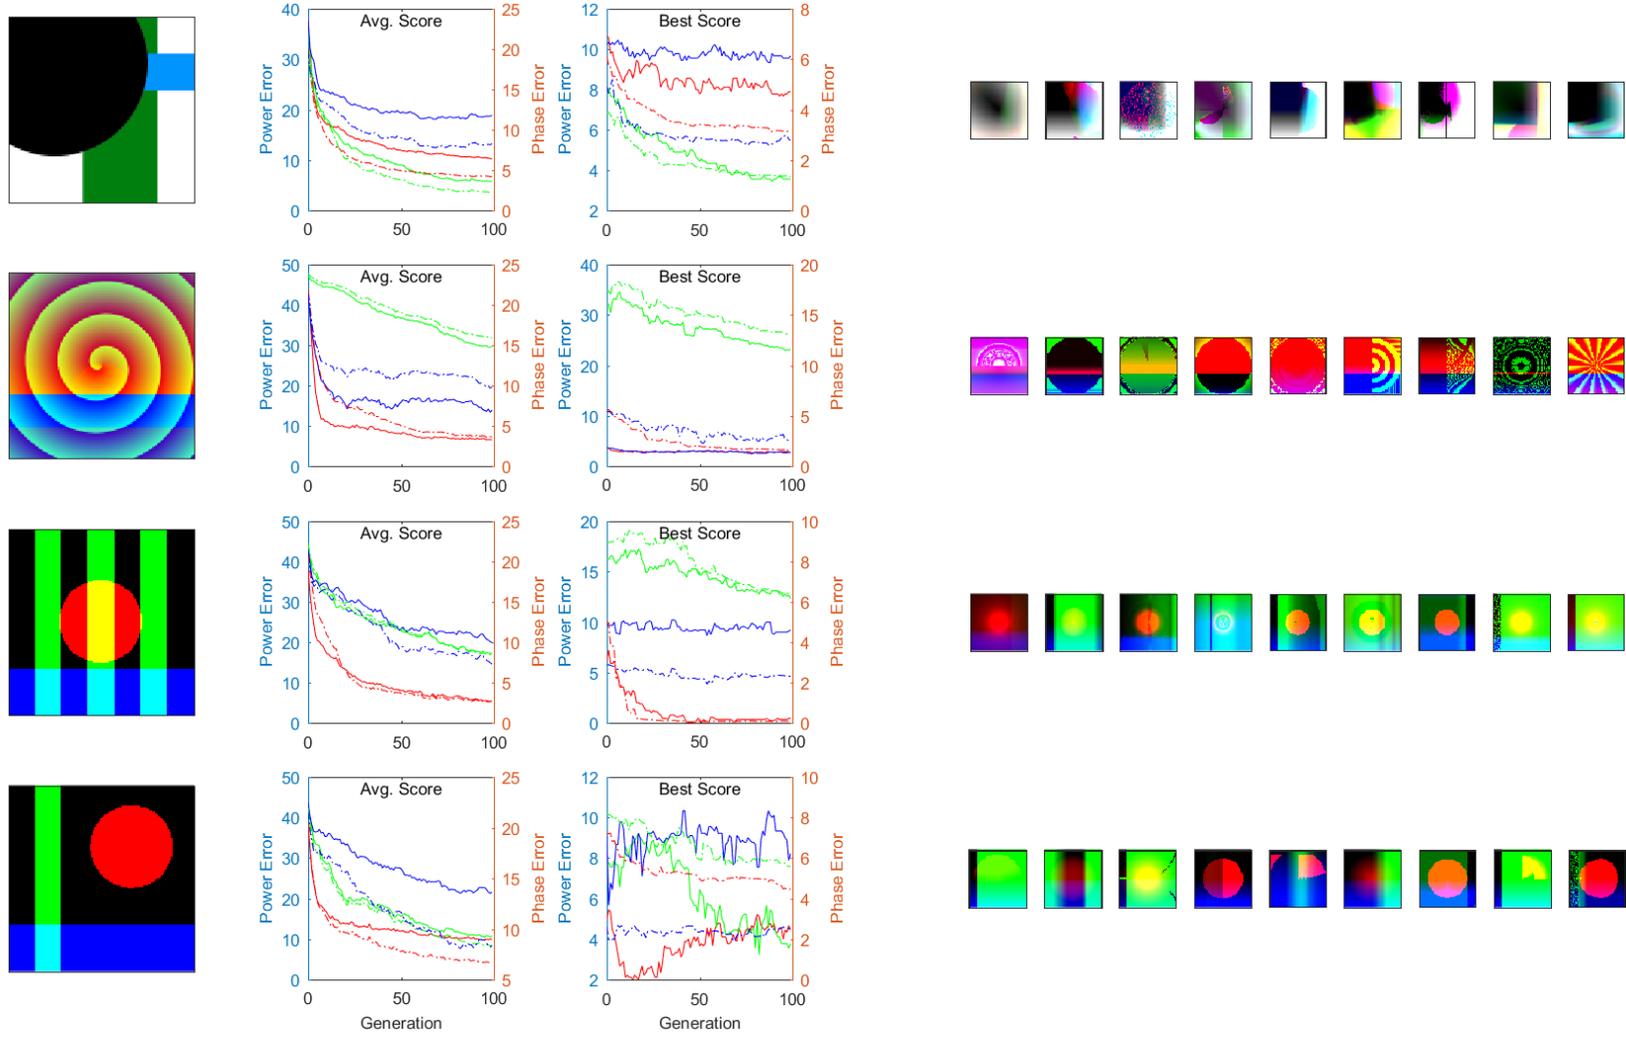
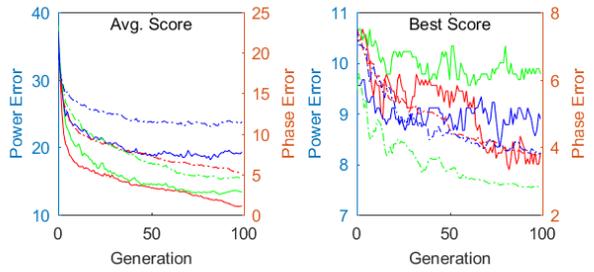
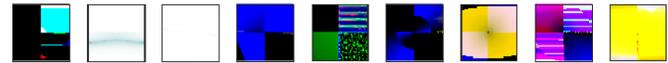
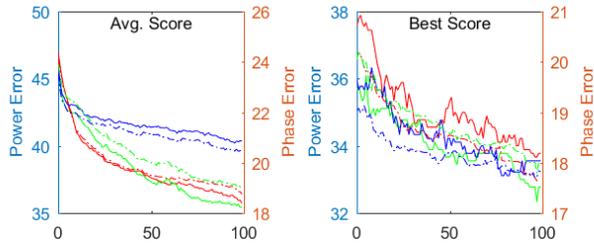
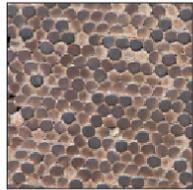
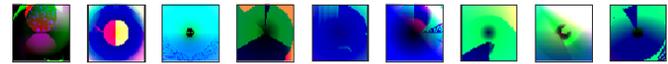
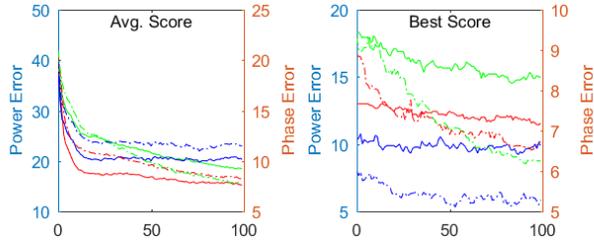
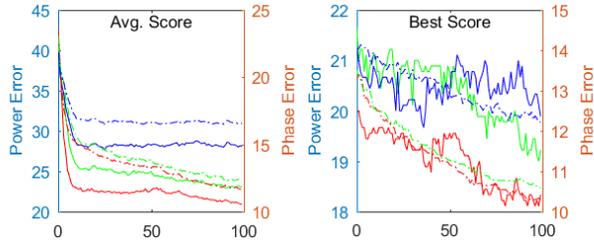
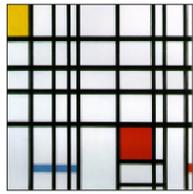


Figure 8.12: Experiment C1 Summary Charts & Examples



sum-of-ranks fitness. The charting for Composition\_15 also shows an increased difficulty in reproducing the red channel phase, which is expected due to the additional offset placed in the red channel of the target. With the Composition\_14 green channel performance being somewhat poor in contrast to its monochrome experiments, we can also conclude that attempting to satisfy an increased number of objectives is increasing difficulty for objectives which could be satisfied in isolation.

When we examine the more complex genre targets, we see results whose spatial features are not dissimilar from the monochrome experiments. These features are not quite optimal, and we had suspected that the targets would be more difficult, but they show similar rough attributes as before. Additionally, we begin to see some colour similarities in a few of the targets. The Mondrian and Cable Ends targets do not show any colour similarities with sufficient consistency to make any conclusions. The Van Gogh, with predominantly blue/cyan shades, has similar colour distributions in its evolved candidates. The flower target also produces candidates with proportions of green and pink/white. This control of colour through relative proportion and overlay of RGB channels, while basic and limited, appears to have success with certain targets.

### Colour Channels: Y+RGB

In an attempt to further constrain the overall composition of the image, and to capture some spatial information lost by considering colour channels in isolation, our next experiment will include additional objectives to guide the power and phase similarity of the average luminance of the evolved colour candidate. It is hoped that by including these additional objectives, attempts to sacrifice any individual colour channel objective will incur further penalties from overall luminance degradation. This scheme will effectively operate on 4 colour channels, requiring a total of 8 objectives in our sum-of-ranks fitness approach.

$$Y = 0.299R + 0.587G + 0.114B \quad (8.1)$$

The conversion from colour (RGB) to grayscale (Y) will be performed using NTSC (CCIR 601) luminance, outlined in Equation 8.1. This provides a close approximation of colorimetric luminance from the non-linear, gamma corrected RGB values. However, due to this weighting, we consider that there might be slight bias toward improving error for the more heavily weighted components.

The inclusion of the additional luminance channel appears to have slightly improved the shape characteristics of the evolved candidates, but we see worsened colour

Figure 8.13: Experiment C2 Summary Charts & 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.

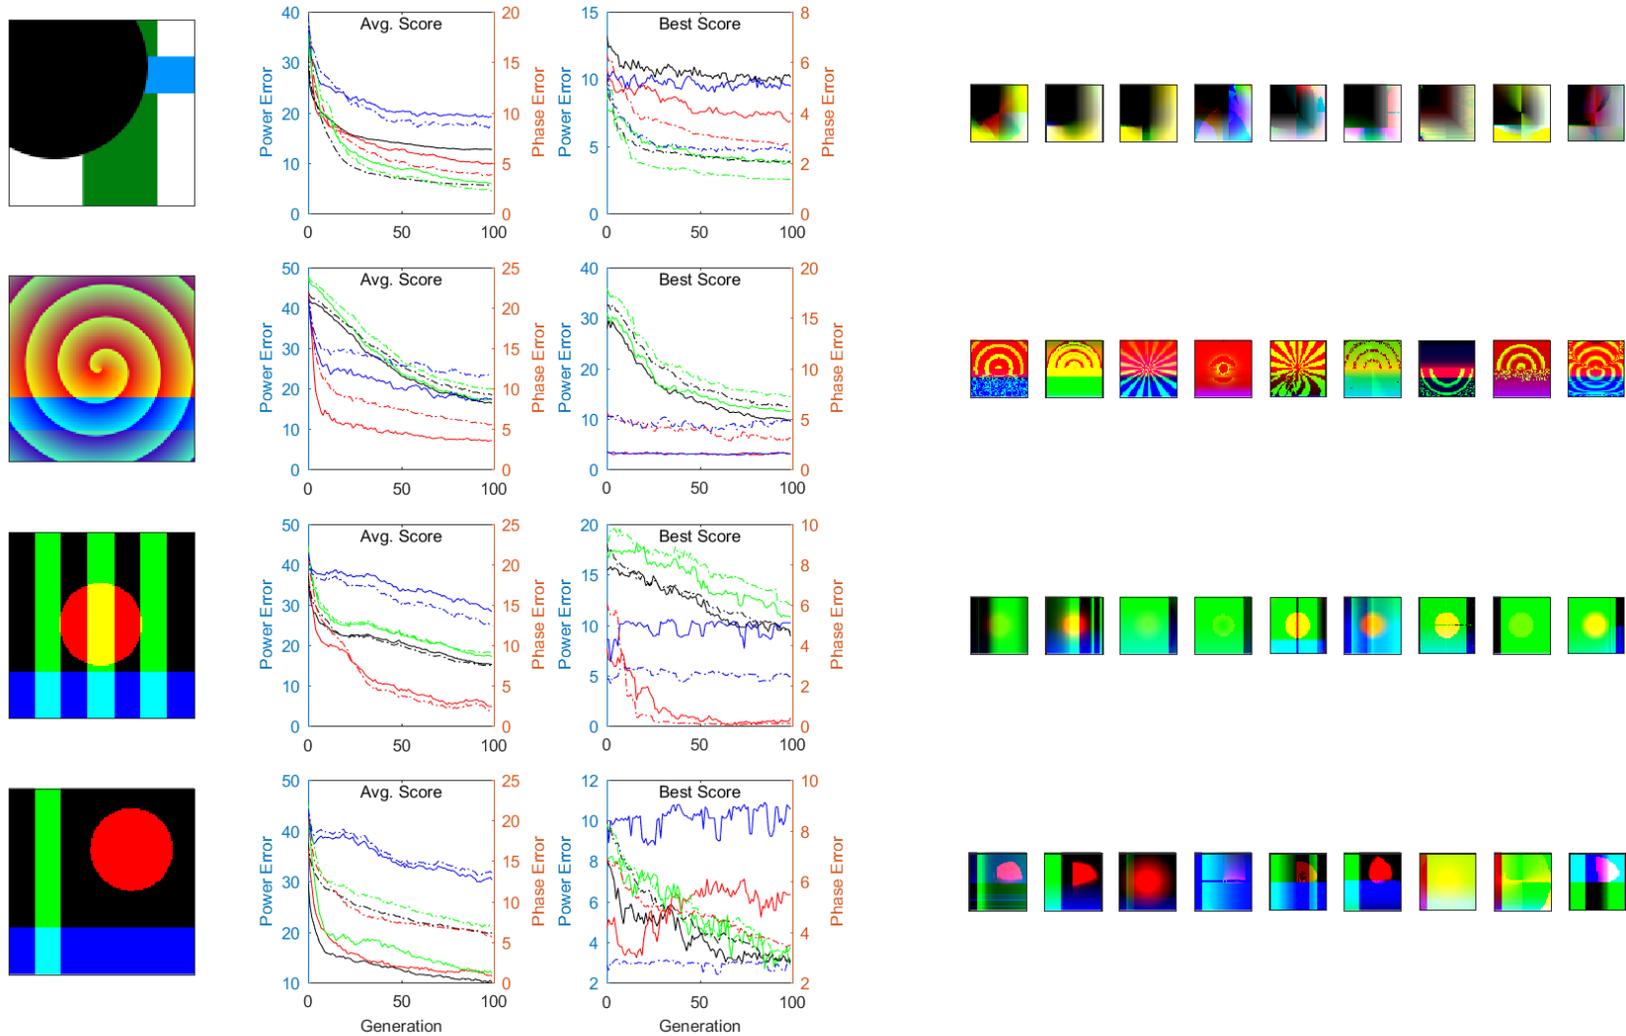
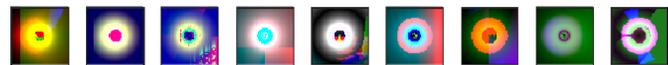
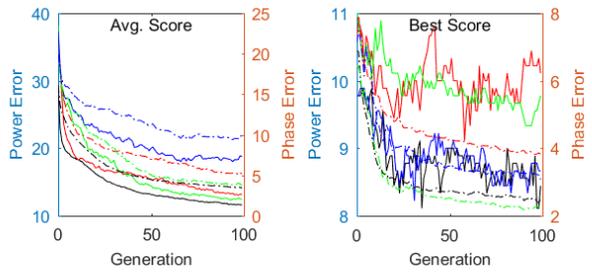
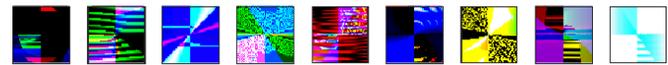
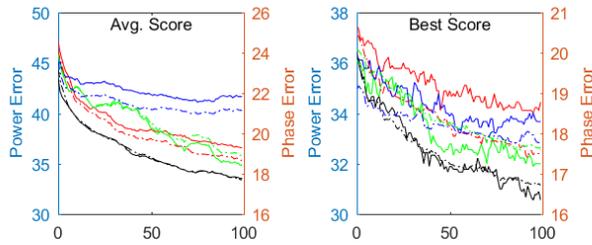
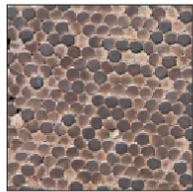
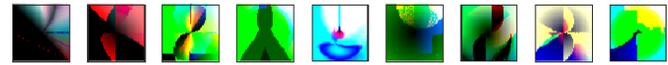
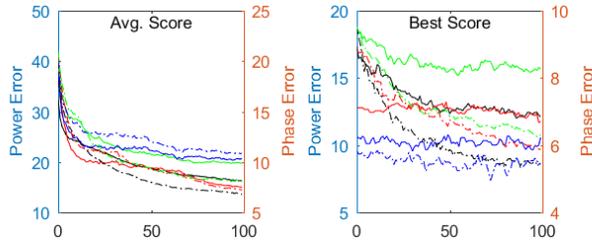
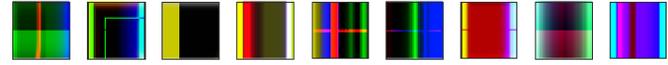
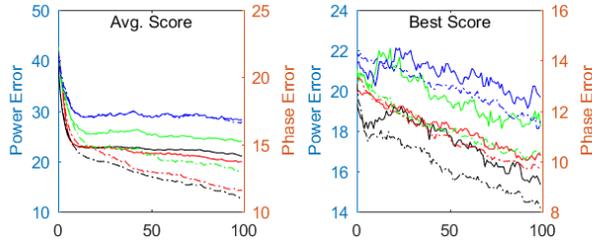
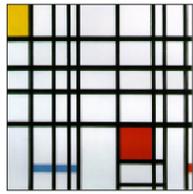


Figure 8.13: Experiment C2 Summary Charts & Examples



similarity across the more complex targets. From the charting in Figure 8.13, 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. We see the Van Gogh and flower targets having marginally improved spatial composition, but all genre targets appear much rougher in colour distribution.

While we had hoped that the inclusion of a luminance channel would reduce the occurrence of sacrifice of 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.

As we had seen improvement in some targets by expanding the GP language, and as our additional luminance channel objectives appear to more heavily constrain shape, we will explore these language changes with the current fitness scheme to see if reduced difficulty if matching shape will improve colour distribution.

In Figure 8.14, we can see the results of including the previously described noise operators. Largely, colour distributions do not appear to have improved. We see our basic compositional targets appear to have visually reduced similarity, though *Composition\_13* is certainly far from lacking creativity. The Van Gogh target appears to have changed, though it is not immediately clear if it is for the worse or better. Notably, we see an improvement to shape with the flower target, including some rather novel specimens. Colour appears to have improved slightly, which is reflected in the terminal scores for the best found individuals per run.

To complete our exploration with language variations in this fitness approach, we lastly reconsider removing the  $X$  coordinate variable from our language, to bias candidates toward using noise operators. We had previously seen some improvements with certain targets using this approach, so we will consider it for completeness.

The unfortunate results of this approach, as seen in Figure 8.15, is that most targets exhibited reduced performance in both shape and colour distribution. The noted exception being the flower target, which, already finding improved candidates using noise, performed similarly to the previous language set. The Mondrian target also saw a marginal improvement to the overall luminance scores, and evolved candidates were more likely to display variations across the vertical axis, though appearances are still very rough approximations to the target.

In an attempt to reduce time constraints from limiting the performance of the

Figure 8.14: Experiment C2 Summary Charts & Examples; Noise Language

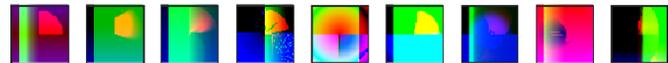
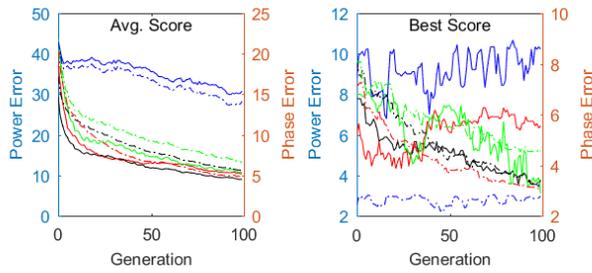
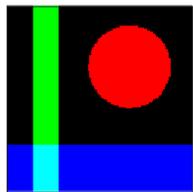
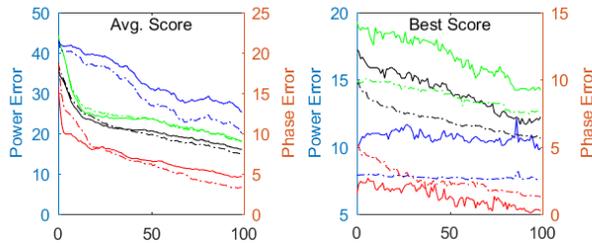
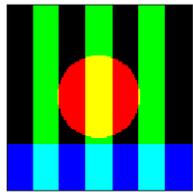
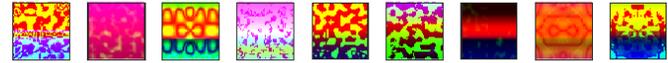
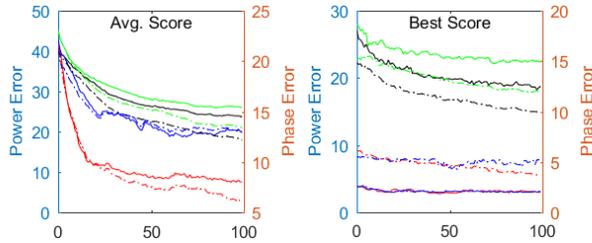
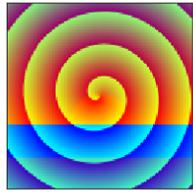
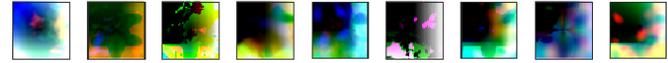
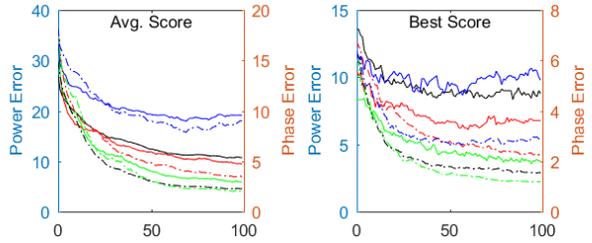
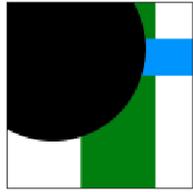


Figure 8.14: Experiment C2 Summary Charts & Examples; Noise Language

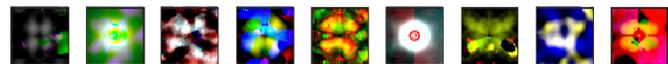
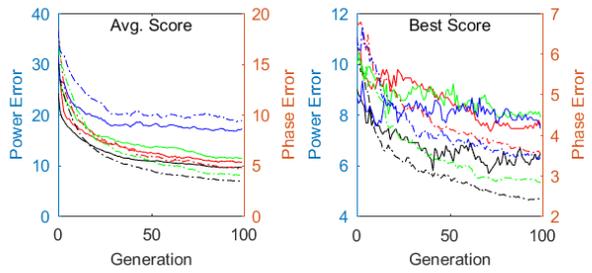
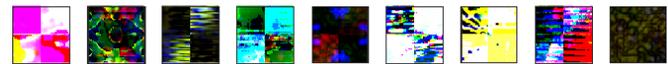
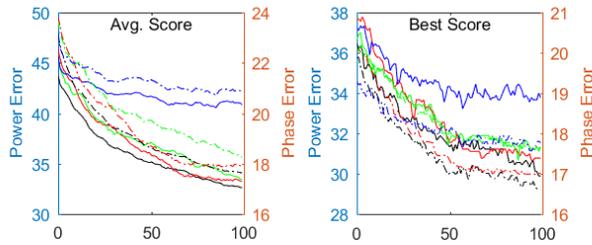
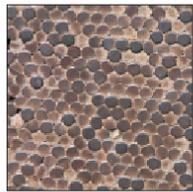
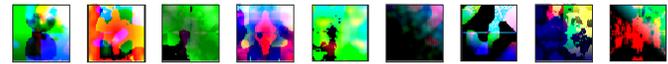
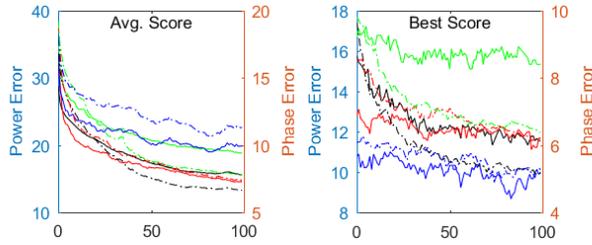
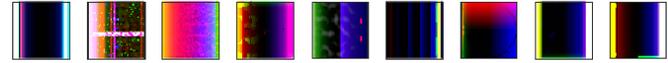
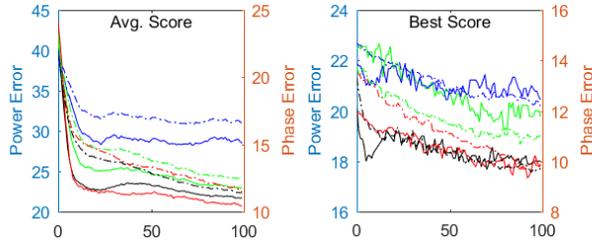
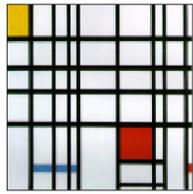


Figure 8.15: Experiment C2 Summary Charts & Examples; Noise, No  $X$

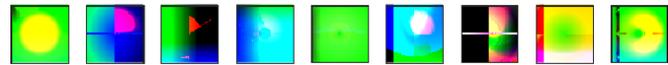
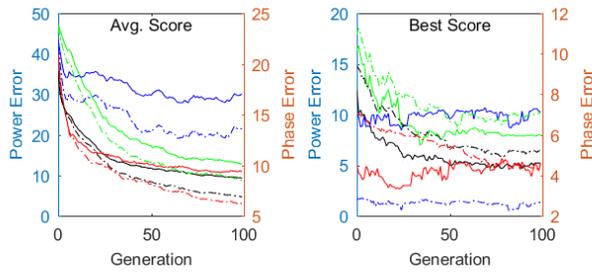
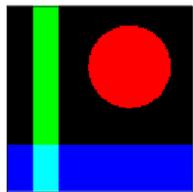
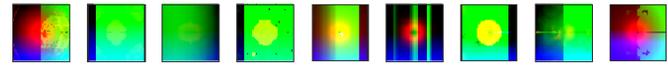
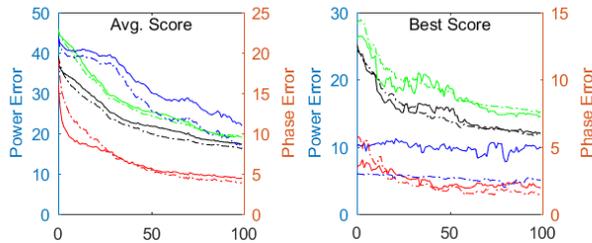
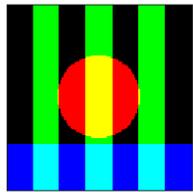
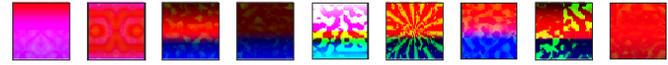
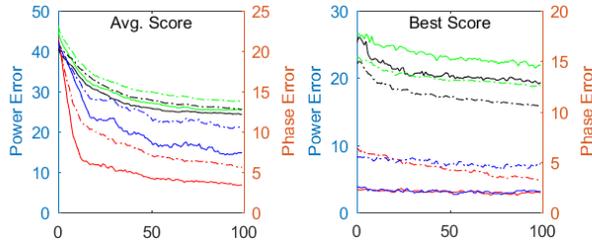
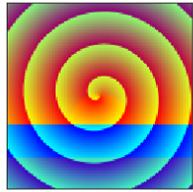
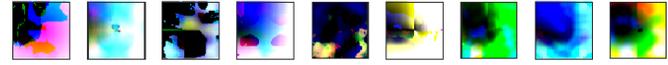
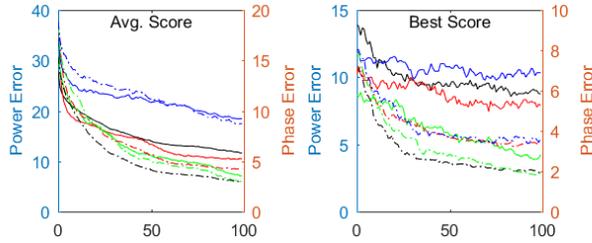
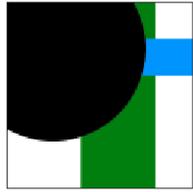
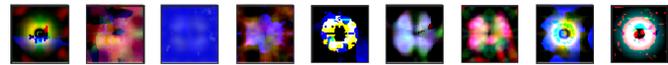
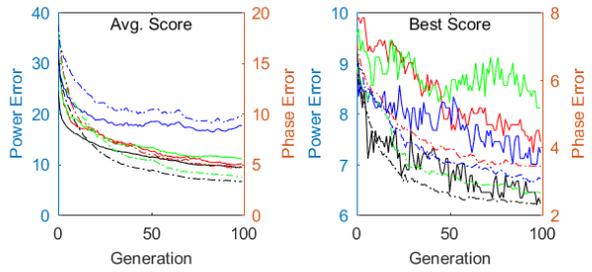
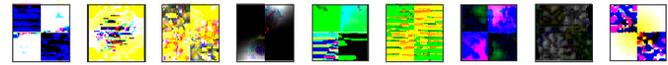
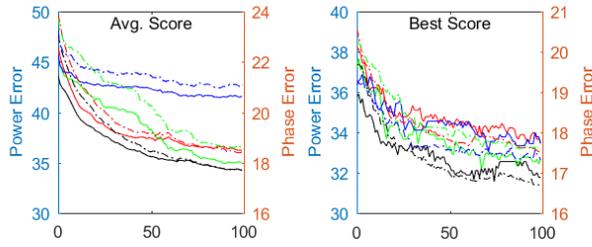
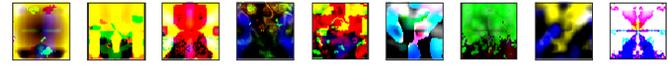
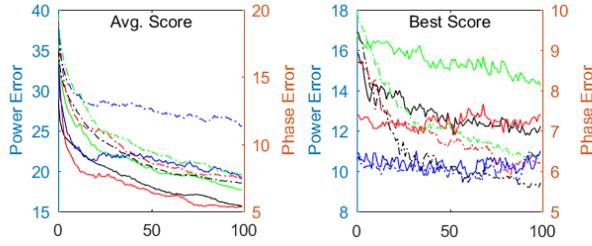
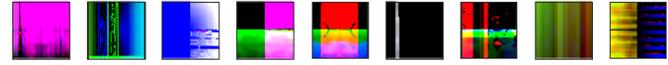
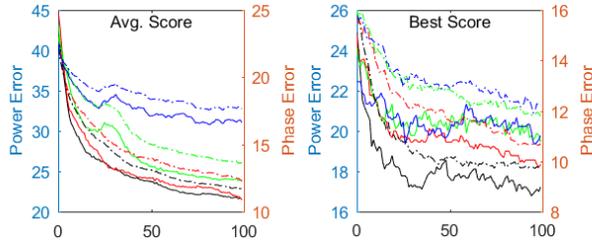
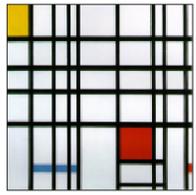


Figure 8.15: Experiment C2 Summary Charts & Examples; Noise, No  $X$



produced results, we have attempted some additional experimentation using an extended GP generation limit of 200. In cursory results, a marginal change to scores was observed where sacrificed objectives were exacerbated and no particular visual changes were noticed. It would seem that candidates had already converged toward particular areas of the search space by that time.

While overall colour distribution could be improved, we did see increased performance when targets held colour channels which could be replicated as grayscale targets individually. As our GP system is constructing its candidates along these same RGB colour channels, this is not an unexpected finding, so much as a matter of sufficiency. While considering these limitations, we were still able to replicate shape and colour for a number of targets, and were able to produce some novel – if rough – variations of the targets. Some highlights have been shown in Figure 8.16.

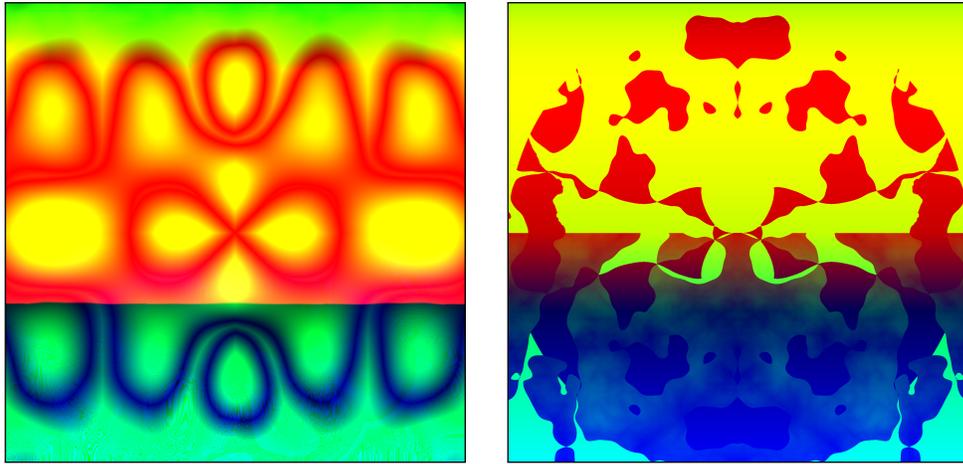
### **Colour Channels: HSL**

From the previous experiment, we saw that using luminance may assist in replicating shape characteristics while guiding evolution in colour images. We also hypothesized and observed that unexpected colours could appear where there were incongruities between candidate and target in even one colour channel. In an attempt to capture changes in perceived colour without these contrasting artifacts, we consider using a different colour channels, such as hue, saturation, and lightness (HSL). Our production of the candidate images will remain the same, reverting back to our basic GP language with polar coordinates added. However, instead of processing the RGB channel values with our power spectra measures, we will first convert our colour channels into the HSL colour space.

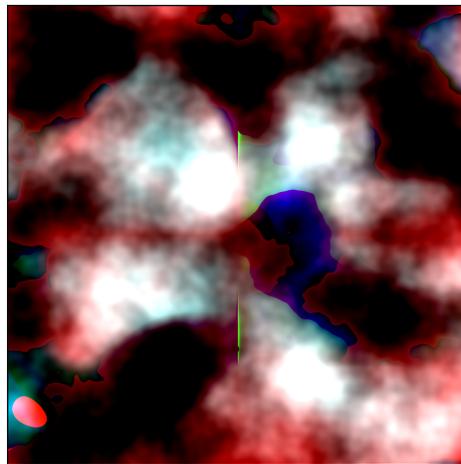
Unfortunately, it would appear that using our fitness measure on HSL channels provides a substantial degradation in visual performance. From Figure 8.17, we can see that many targets have poor spatial similarity, and almost all fail in obtaining matching colour distributions. Many of the targets which perform well numerically still show very rough visual performance. Some interesting observations can be made from *Composition\_13*, *Composition\_15*, and the *flower* target. For these targets, we see some reasonable, though rough, spatial similarities, and comparatively low numerical error. Notable, is that these shapes emerge without being constrained to the original RGB channels of the target, and that the colours of these shapes are actually quite different. We would expect alignment of the correct hue to be determined by hue phase measures. Shapes emerging with incorrect hue may then suggest that further pressure and accuracy will be required with regards to phase in this channel, and

Figure 8.16: Experiment C2 Highlights

(a) Experiment C2; Composition\_13; Noise language



(b) Experiment C2; Flower; Noise language



(c) Experiment C2; Flower; Noise language, no  $X$

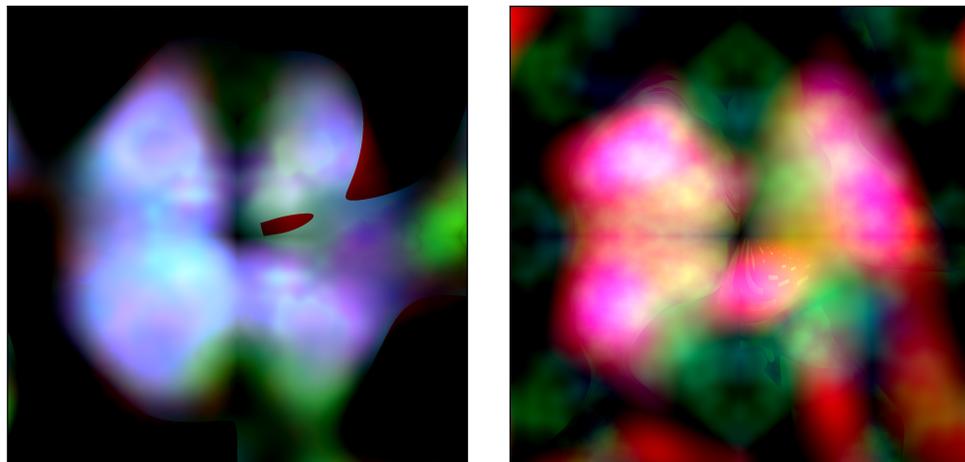


Figure 8.17: Experiment C3 Summary Charts & Examples

Power and phase errors for each of the hue, saturation, and luminance colour channels are plotted using cyan, magenta, and black respectively. Power error is denoted with solid lines, where phase error uses a dashed line.

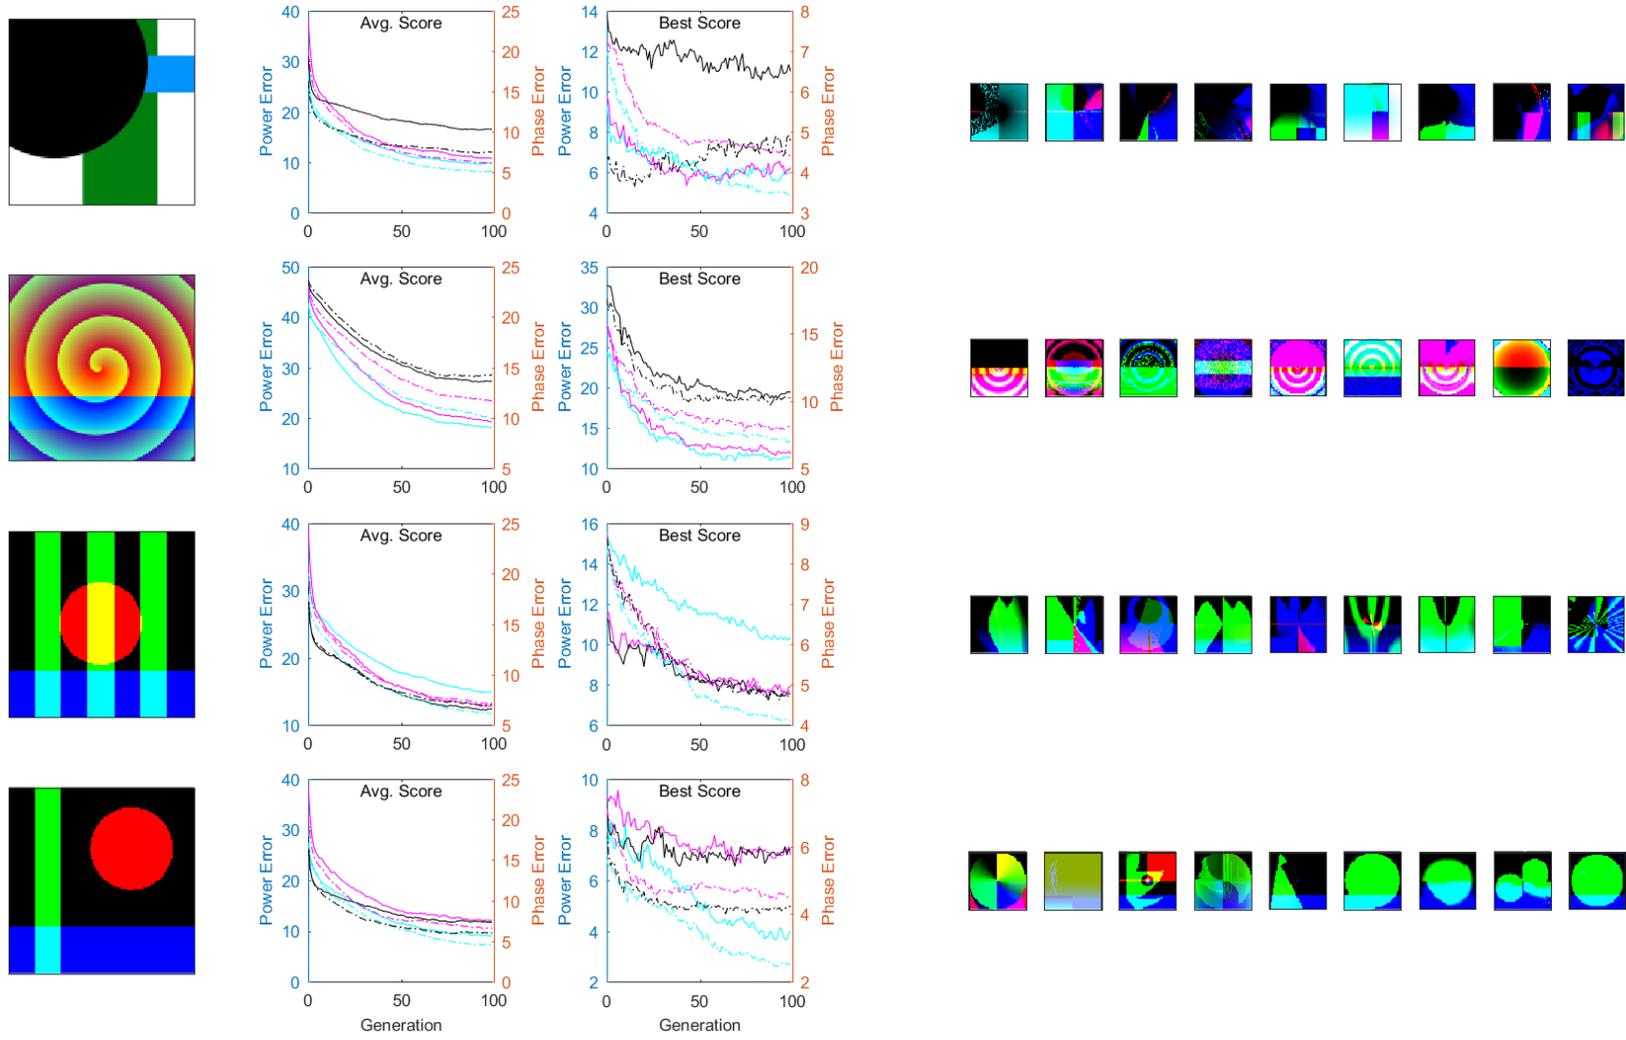
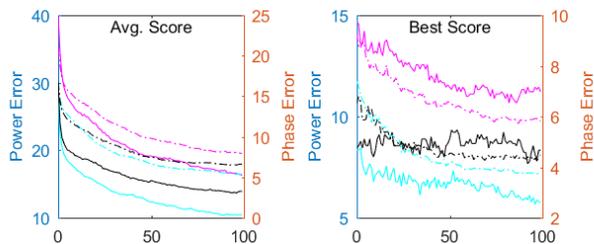
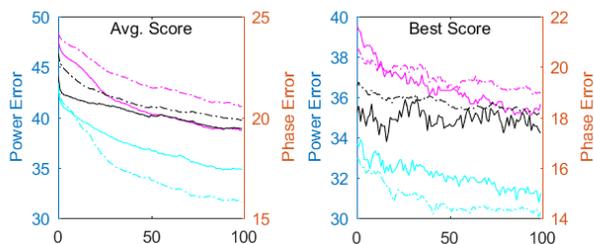
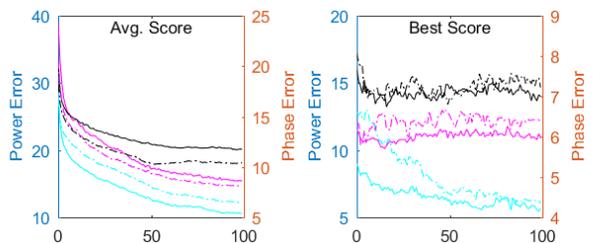
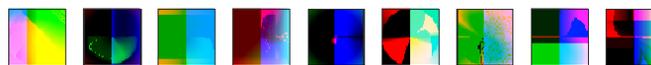
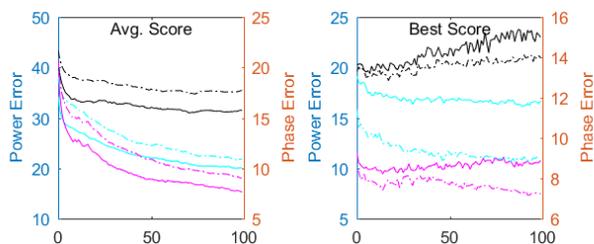
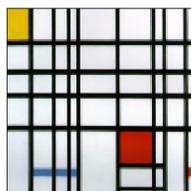


Figure 8.17: Experiment C3 Summary Charts & Examples



that an adjustment to the P1 fitness scheme may be needed if we are to pursue this approach to colour.

Potentially, the disparity between constructing images across RGB channels and evaluating them using HSL could be creating an extra amount of difficulty for the system when guiding image evolution. Small changes to a single GP tree, while only changing a single RGB colour channel, may be creating wide-scale changes to the evaluated HSL representations. Reproductive operations on the GP trees may then be seen as overly destructive, and the system shows difficulty in guiding candidates.

### 8.2.3 Separating Shape & Colour

We see success with certain targets when using the previous fitness measures on the isolated colour channels, but there is an obvious alternate approach to guiding similar colour distributions: using separate colour and shape objectives. To capture shape, we will use the existing P1 fitness measure across the grayscale luminance channel as discussed in experiment C2. To capture and guide colour, we will consider the CHISTQ measure as proposed by QBIC Project [57].

There are numerous measures for colour similarity which have seen use in search and retrieval systems, though the CHISTQ measure, a histogram difference weighted by colour similarity, has also previously shown promise with evolutionary texture systems [44][70]. The CHISTQ measure first finds the differences between each of the target and candidate histogram bins. Each bin difference is then factored with every other bin difference and a similarity measure, all summed together to form the histogram distance weighted by colour similarity. The goal is to provide a measure of histogram similarity which is able to account for perceptual distance between differences in pairs of quantized colours. This is more formally defined in Equation 8.2 [57].

$$CHISTQ = Z^T AZ \quad (8.2)$$

$$A(i, j) = 1 - d(c_i, c_j)/d_{max} \quad (8.3)$$

After quantizing both target and candidate images into suitable bins, we can produce the normalized element by element difference (or bin difference)  $Z$ . We will have  $A$  denote a normalized, symmetric colour similarity matrix, as defined by Equation 8.3. Results from the QBIC Project found MTM and LUV colour distance to perform well for their retrieval system, though we will adopt a simple euclidean distance in RGB

colour space,  $d$ , as used by [70]. We will also follow with their choice of binning to 3 bits per colour channel.

In experiments using both the polar coordinate language set, and its expansion to noise operators, the C4 experiments using the CHISTQ objective produce results with spatial attributes similar to their C2 experiment counterparts. In Figures 8.18 and 8.19, we see familiar compositional reproductions, and similar properties with the replicated genre targets. When we render the evolved candidates into a grayscale representation, we find that there are the same style of spatial properties, and similar performance for the genre targets as when evolving grayscale candidates.

Composition\_13-15 have candidates which show some colour properties common to their targets, and proportional lower error scores in the CHISTQ objective. Many of the other targets have dissimilar colour distributions and higher CHISTQ errors. It appears that, while the CHISTQ objective converges when measured through the population mean, the best found individuals routinely ignore the measure to accommodate the other two semi-dependant spatial objectives. It is possible that for Composition\_13-15, produced with simpler compositions in each colour channel, guiding the evolutionary search space toward spatial matches also worked favourably for colour similarity. Conversely, the more complex targets would need further incentive to yield spatial traits for colour, possibly considering adjusted weighting between sum-of-ranks objectives.

As an interesting aside, we found that the CHISTQ measure also had a notable influence on the more complex, grayscale compositional targets. In Figure 8.20, we note that the addition of the CHISTQ objective has led to grayscale candidates which have a more similar proportion of intensity levels, with only minor loss of spatial similarity.

Figure 8.18: Experiment C4 Summary Charts & Examples

Keeping in-line with the previous recent figures, power and phase errors for across the average luminance have been plotted in black. Power error is denoted with solid lines, where phase error uses a dashed line. The CHISTQ error measure is plotted in magenta along the phase axes, scaled by a power of 10 to maintain a reasonable graph window. Grayscale interpretations of evolved candidates are shown beneath their colour renditions.

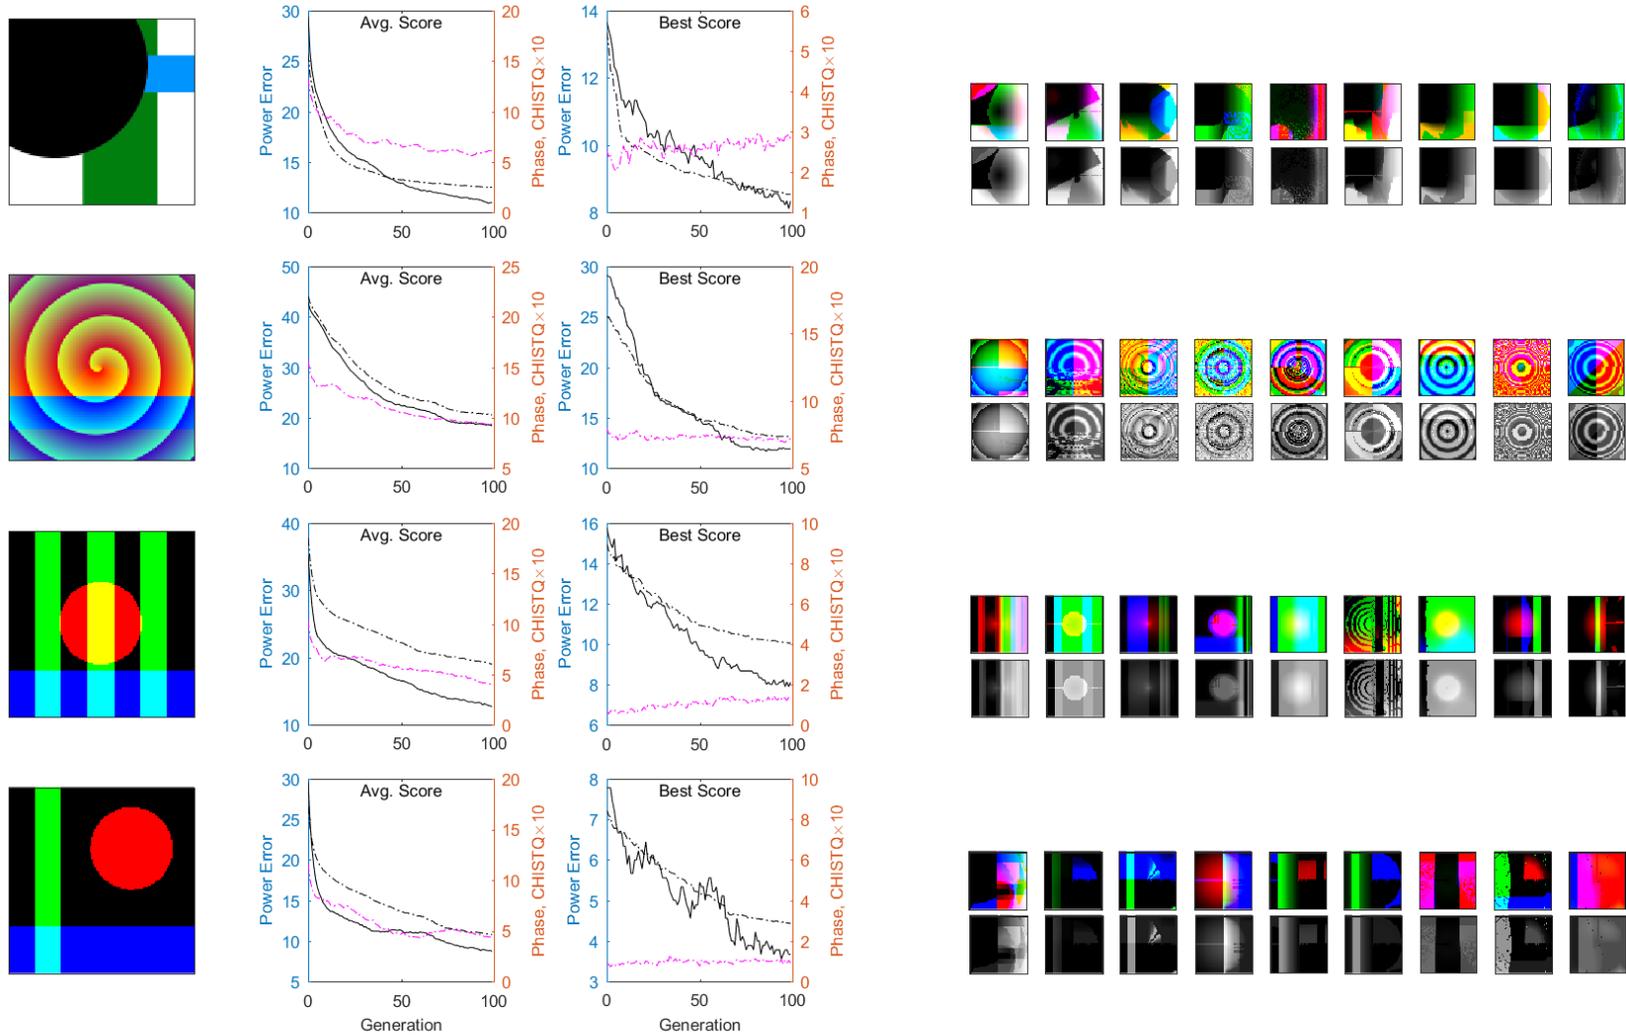


Figure 8.18: Experiment C4 Summary Charts & Examples

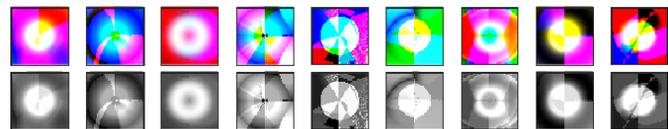
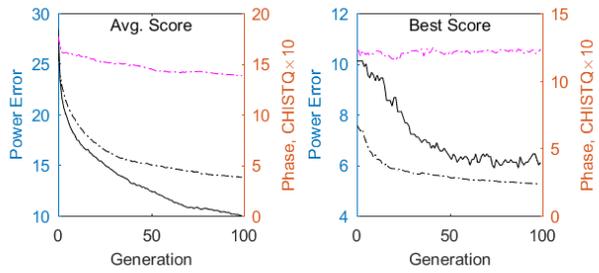
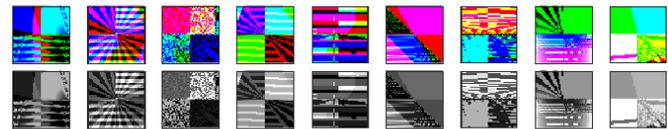
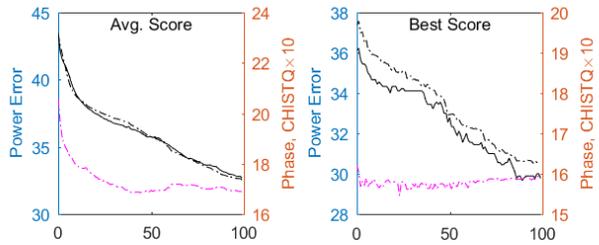
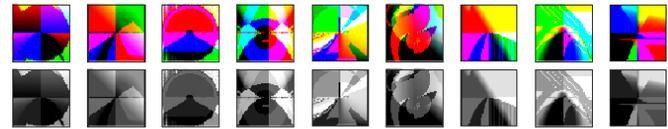
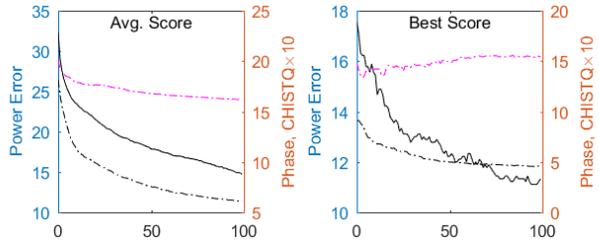
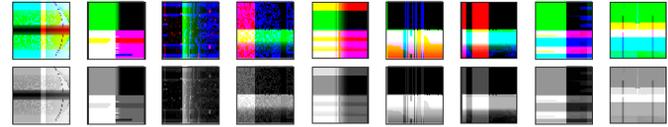
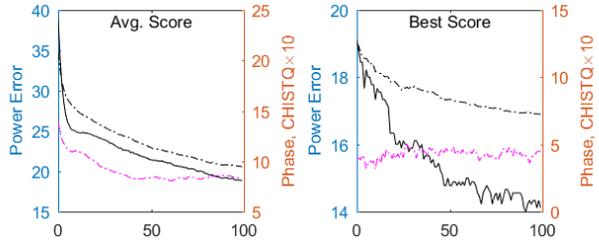
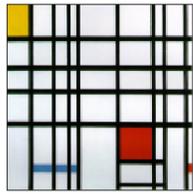


Figure 8.19: Experiment C4 Summary Charts & Examples; Noise Language

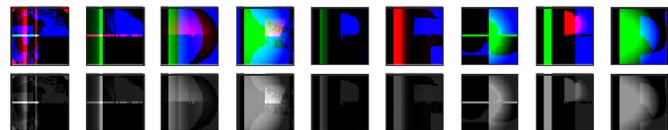
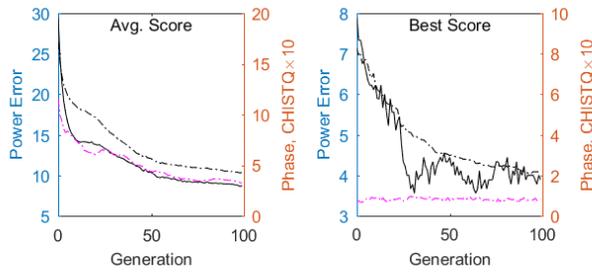
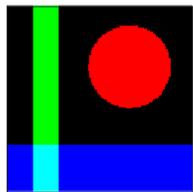
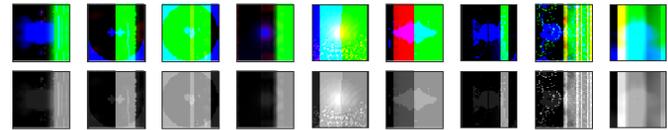
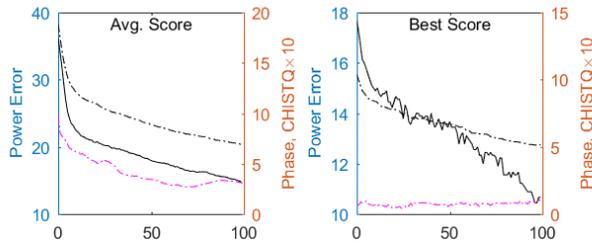
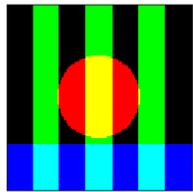
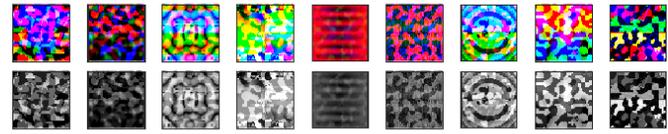
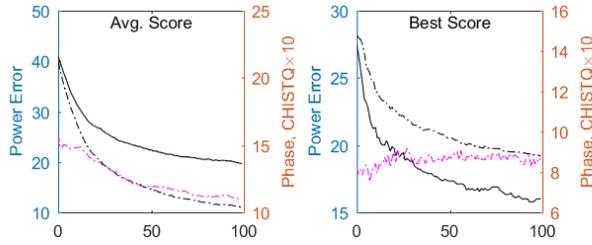
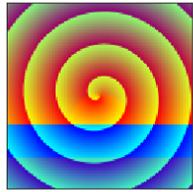
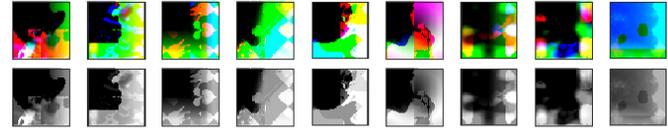
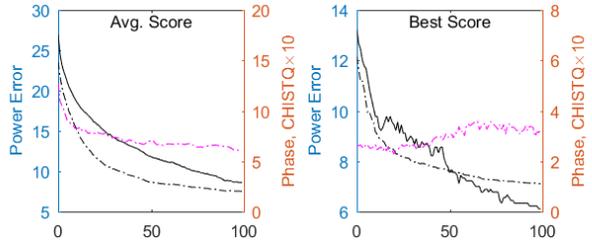
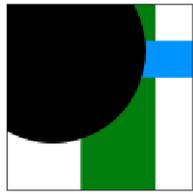


Figure 8.19: Experiment C4 Summary Charts & Examples; Noise Language

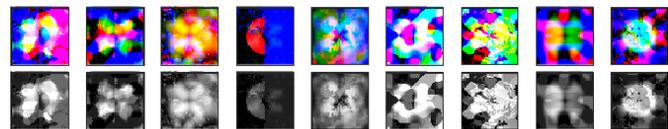
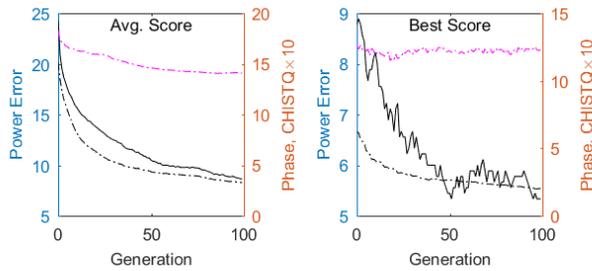
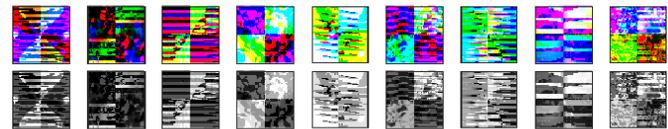
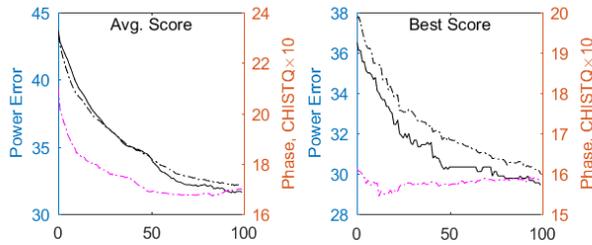
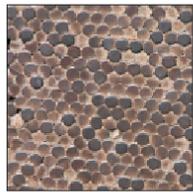
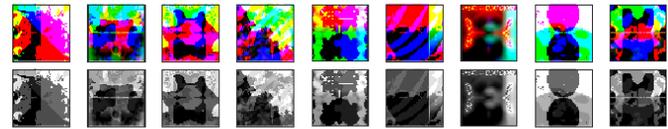
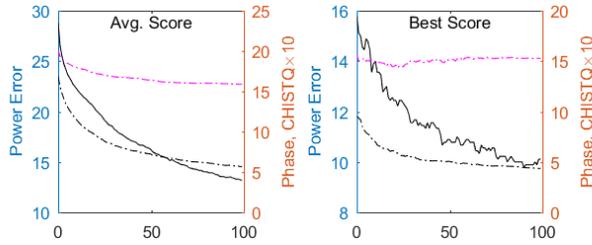
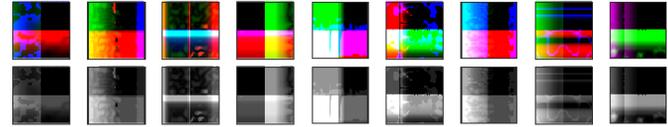
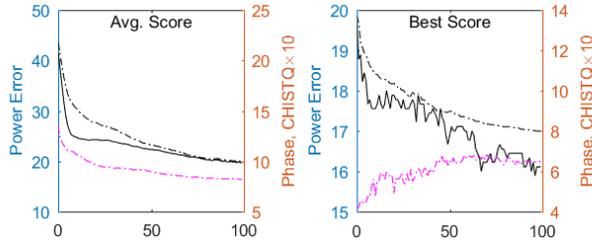
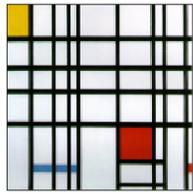
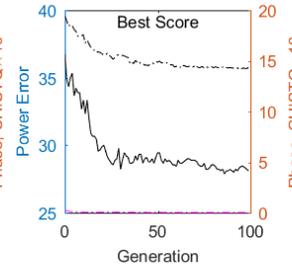
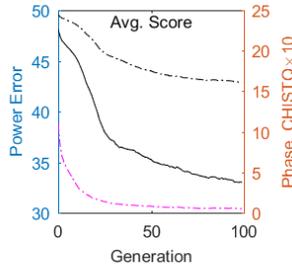
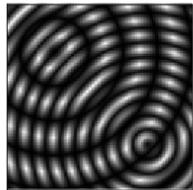
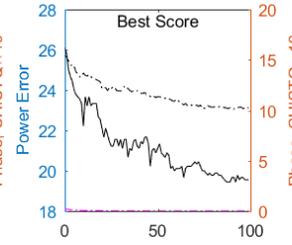
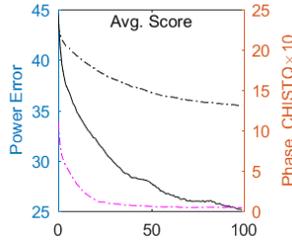
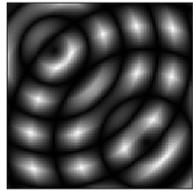
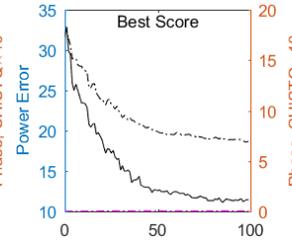
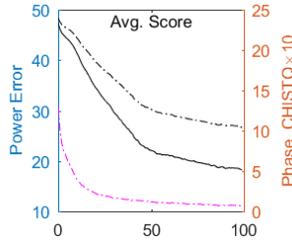
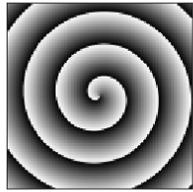
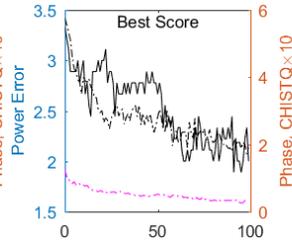
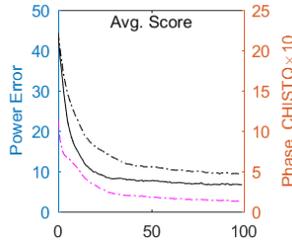


Figure 8.20: Experiment C4 Grayscale Compositional Summary Charts & Examples

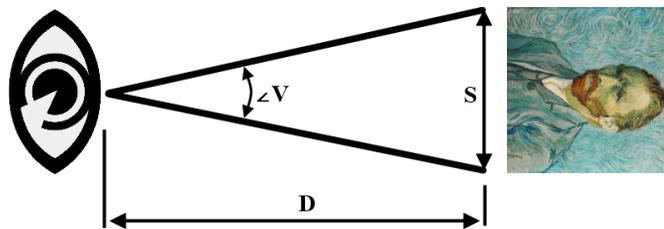


### 8.3 Spatial Frequencies & Comfort

In the course of evolving the many candidate images with each target and experiment set, we identified a number of evolved images which we believe displayed unpleasant or uncomfortable spatial aspects. Previous research from Fernandez and Wilkins [41] found correlations between intensity level contrasts at certain spatial frequencies with increased levels of discomfort.

To further discuss the findings of Fernandez and Wilkins, we need to adjust some of our terminology. The concept of spatial frequency denotes a cyclical nature across a measured space, such as the re-occurrence 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 have 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 of visual angle – when paired with known viewing distance and image size – to compute a relative measure of angular spatial frequency.

Figure 8.21: Calculation of Spatial Frequency

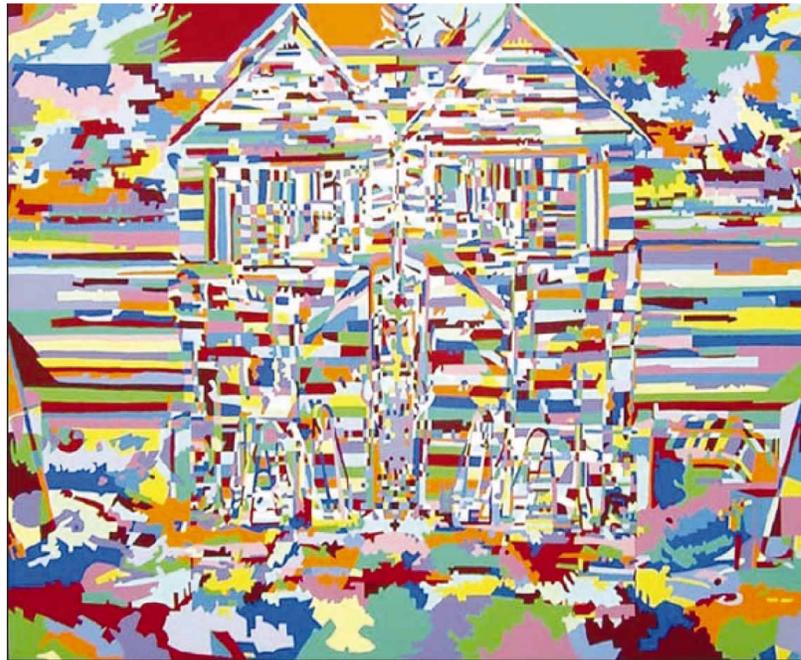


$$V = 2 \arctan \left( \frac{S}{2D} \right) \quad (8.4)$$

With spatial frequencies known in terms of fractional image width, we can use Equation 8.4 to interpolate their corresponding visual angle  $V$  when observed with size  $S$  at a distance of  $D$ .

Returning to Fernandez and Wilkins, their study observed that images with increased amplitudes at a few octaves around 3 cycles per visual degree corresponded with higher reports of image discomfort. We have explored numerous schemes in the previous sections to constrain and obtain specific spatial frequencies of a target image in our newly evolved candidates. With a direct relation between relative visual degree and absolute image spatial frequencies, we hold high hopes that these findings may be combined to the effect of a new aesthetic model.

Figure 8.22: Jesmond Barn by Debbie Ayles © 2003, with inspiration from a basilar artery migraine. Additional images can be found at the artist's website: <http://www.debbieaylesartist.co.uk/varicolour>.

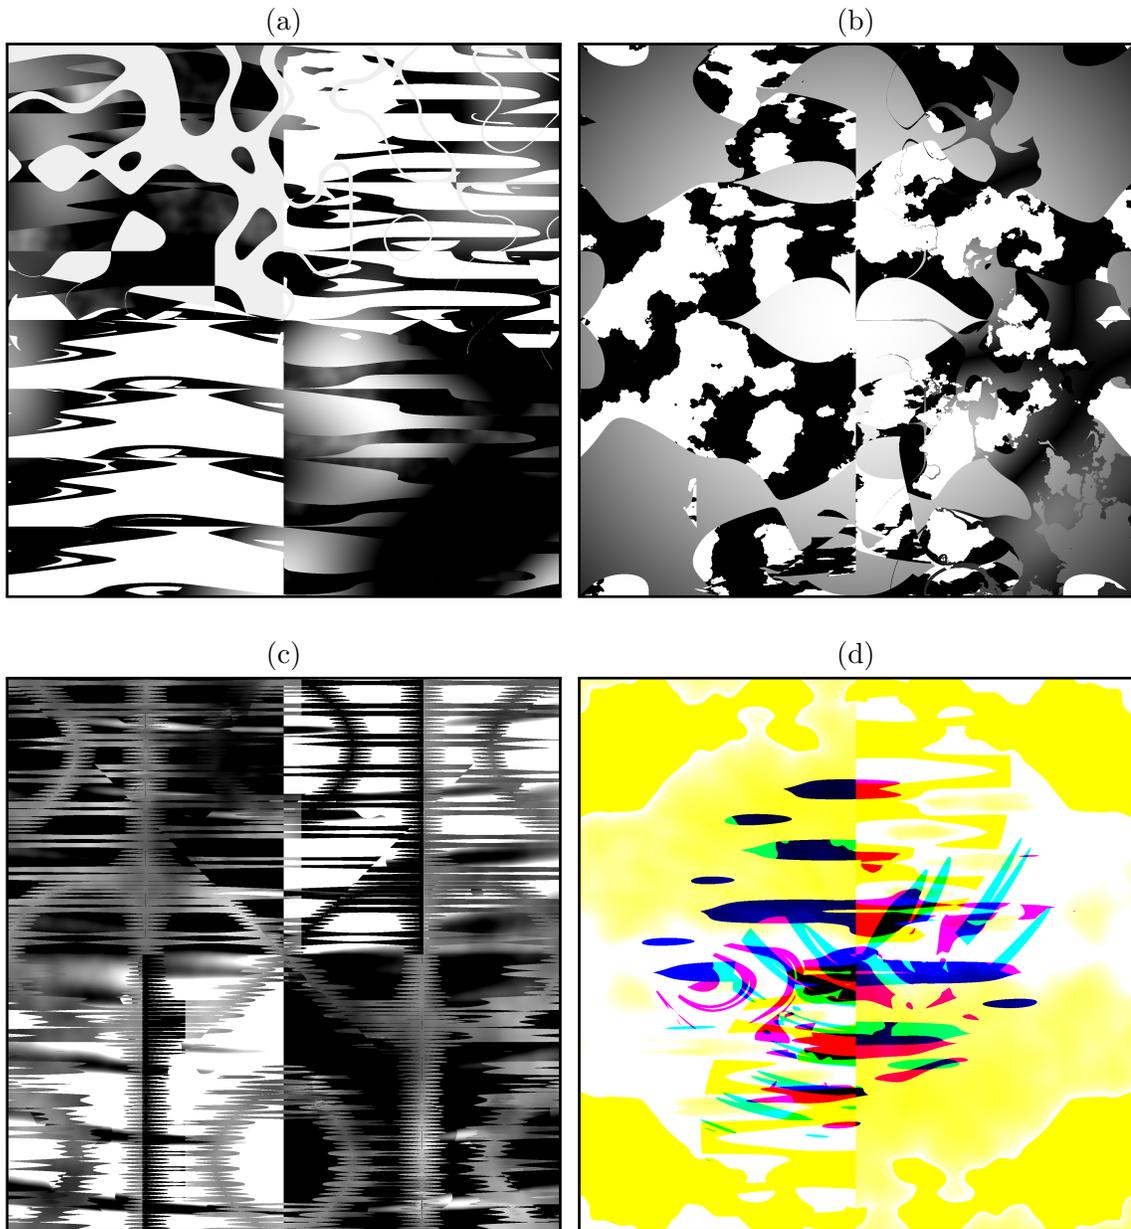


If we are to accept that certain angular spatial frequencies can produce discomfort, then it should follow that the same image can be more or less pleasant depending on its size or distance when viewed. To this end, we can see a number of example images which appear to display these properties. One interesting example is seen in Figure 8.22, where an observation made at a distance may more easily reveal details than when examined up close (though these details once again blur once some distance threshold is exceeded). From afar, we may more readily observe the creativity of the piece and can identify parts of the painted structure. Conversely, we have received a number of anecdotes suggesting visual strain or discomfort when viewed more closely. This phenomenon is by no means limited to this individual piece. Some images we have found in the course of our experimentation which were noted for displaying uncomfortable aspects have been highlighted in Figure 8.23 (resized to fit in this report and to reduce discomfort).

It has been theorized (in *spatial frequency theory*) that the human visual cortex operates through analysis of light receptor spatial frequencies [38][39]. With supporting works finding sensitivity in animals to certain spatial frequency ranges [54][71], it is not surprising to think that humans may also be more sensitive to contrasts at certain spatial frequencies. As seen through the previous image examples, we can

Figure 8.23: 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.)



corroborate that intensity contrasts at certain visual frequencies are uncomfortable; seen as discordant and at times even painful.

As angular spatial frequencies require a defined image size and distance, we should consider how these could affect our analysis. We have used a consistent image size of 12", at a viewing distances of 24"-32", which was chosen to mirror our office configuration. With increased image size producing the inverse effects of increased viewing distance, we can extrapolate results and limit our examination to changing the latter. In Figure 8.24, we examine an image that was immediately noted for uncomfortable aspects, and explore how small changes in view distance can offset which angular spatial frequencies have high power. As we might expect, an increase in view distance (or decrease in image size) shifts power to the lower angular spatial frequencies, and a decrease in view distance (or increase in image size) aligns power with the higher angular spatial frequencies.

At our initial observation distances of 24"-32", the noted mark for 3.0 cycles/degree aligns closely to a pair of more highly powered frequency ranges. Minor changes to viewing distance leave these elevated ranges as still firmly inside the " $3.0 \pm 2$  octaves" suggestion. When we alter the viewing distance to 50", and no longer observe the same visual strain as before, we unfortunately still note the presence of the highly powered frequency ranges within an octave of 3.0 cycles/degree. Conceivably, the image may still be displaying less desirable angular spatial frequencies, though not to the extent of noted discomfort. Nonetheless, it may be necessary to further refine the prioritized frequency ranges.

When we consider some additional image samples noted for their discomfort, we at times find a lack of elevated frequency ranges shown through the radial average, such as the examples in Figure 8.25. We were also surprised to find that our prototypical example, extolled by Fernandez and Wilkins' paper [41], shows a similar lack of elevated power around their proposed frequency ranges (Figure 8.26). Noting the same concerns of excessive abstraction that were made during our naïve fitness schemes, we consider an alternate approach to the radial average which might better capture the observed contrasts in the image. When we extract the maximum error found amongst the radial bins, instead of the average, we can find some slightly stronger frequencies near the targeted 3.0 cycles/degree, though it is not as clear if those are indeed the responsible frequencies.

With these findings, a couple of foreseen limitations can be considered. The first, and least negotiable concern, holds that viewing size and distance must be considered before evolution. While minor adjustments can be made without inducing sudden un-

Figure 8.24: Angular Spatial Frequency Analysis - Distance Variations

At the top left we have the analysed image, beneath which is their 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 are the spatial frequencies calculated relative to visual angle at various viewing distances. As recommended by Fernandez and Wilkins, octaves about 3 cycles/degree have been marked.

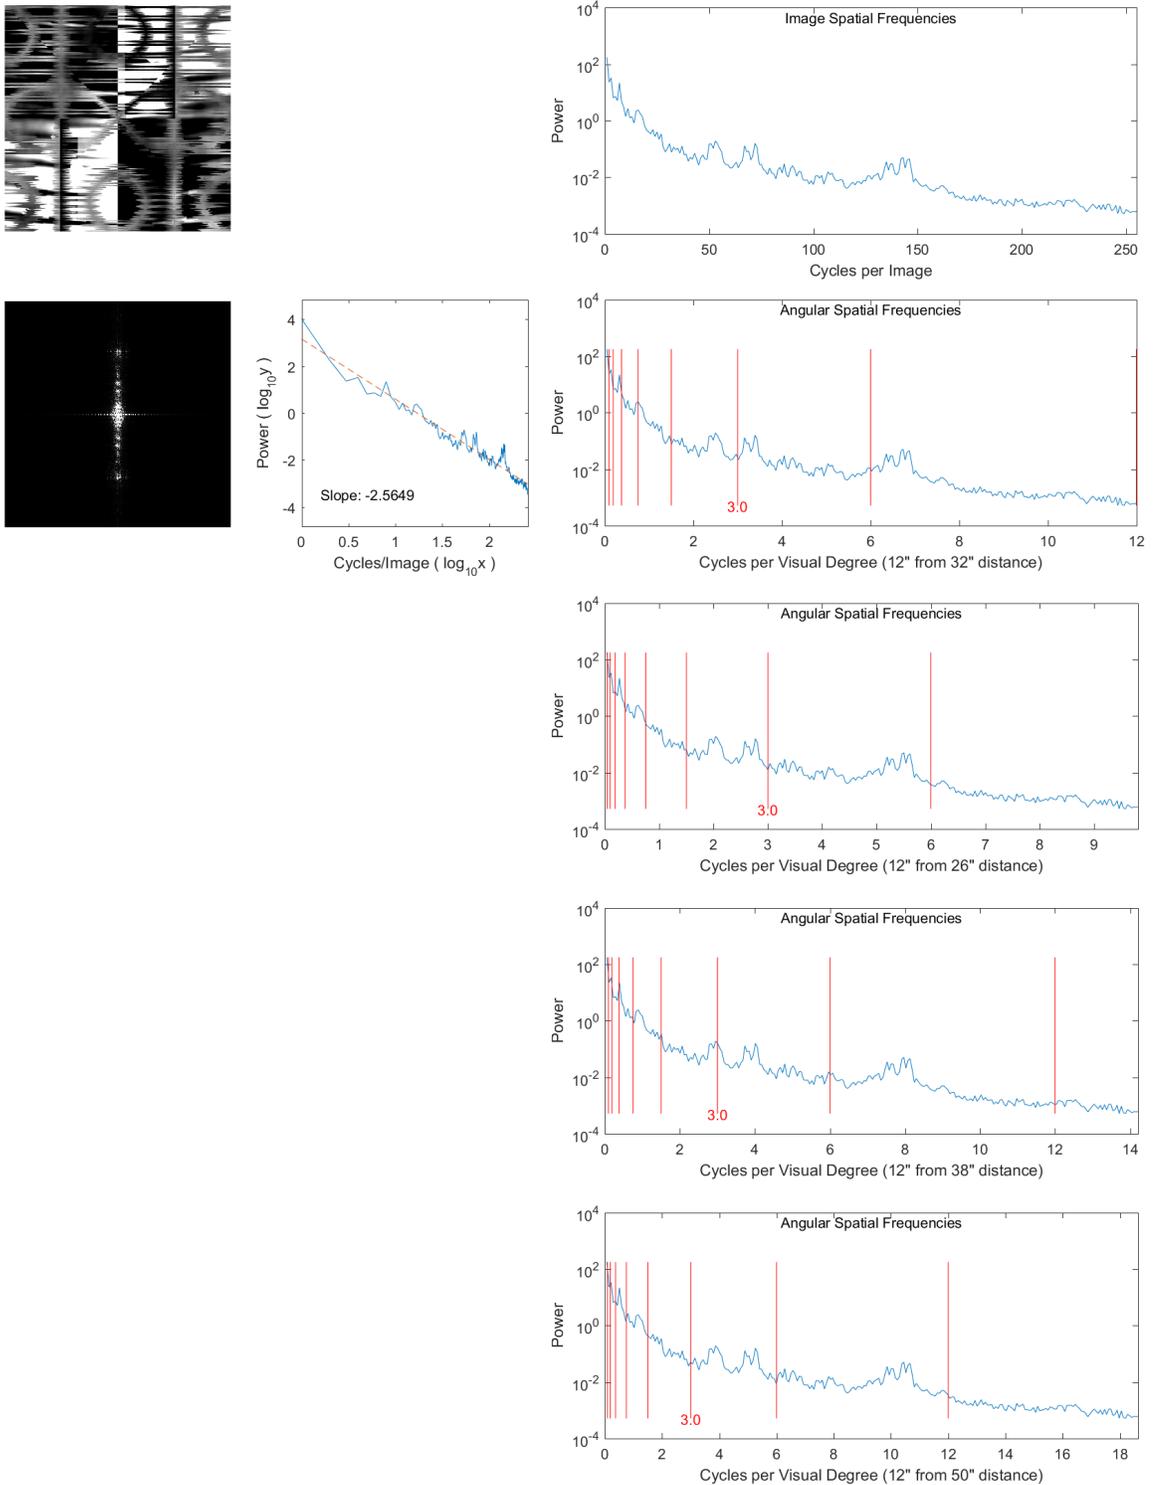


Figure 8.25: Angular Spatial Frequency Analysis: Various

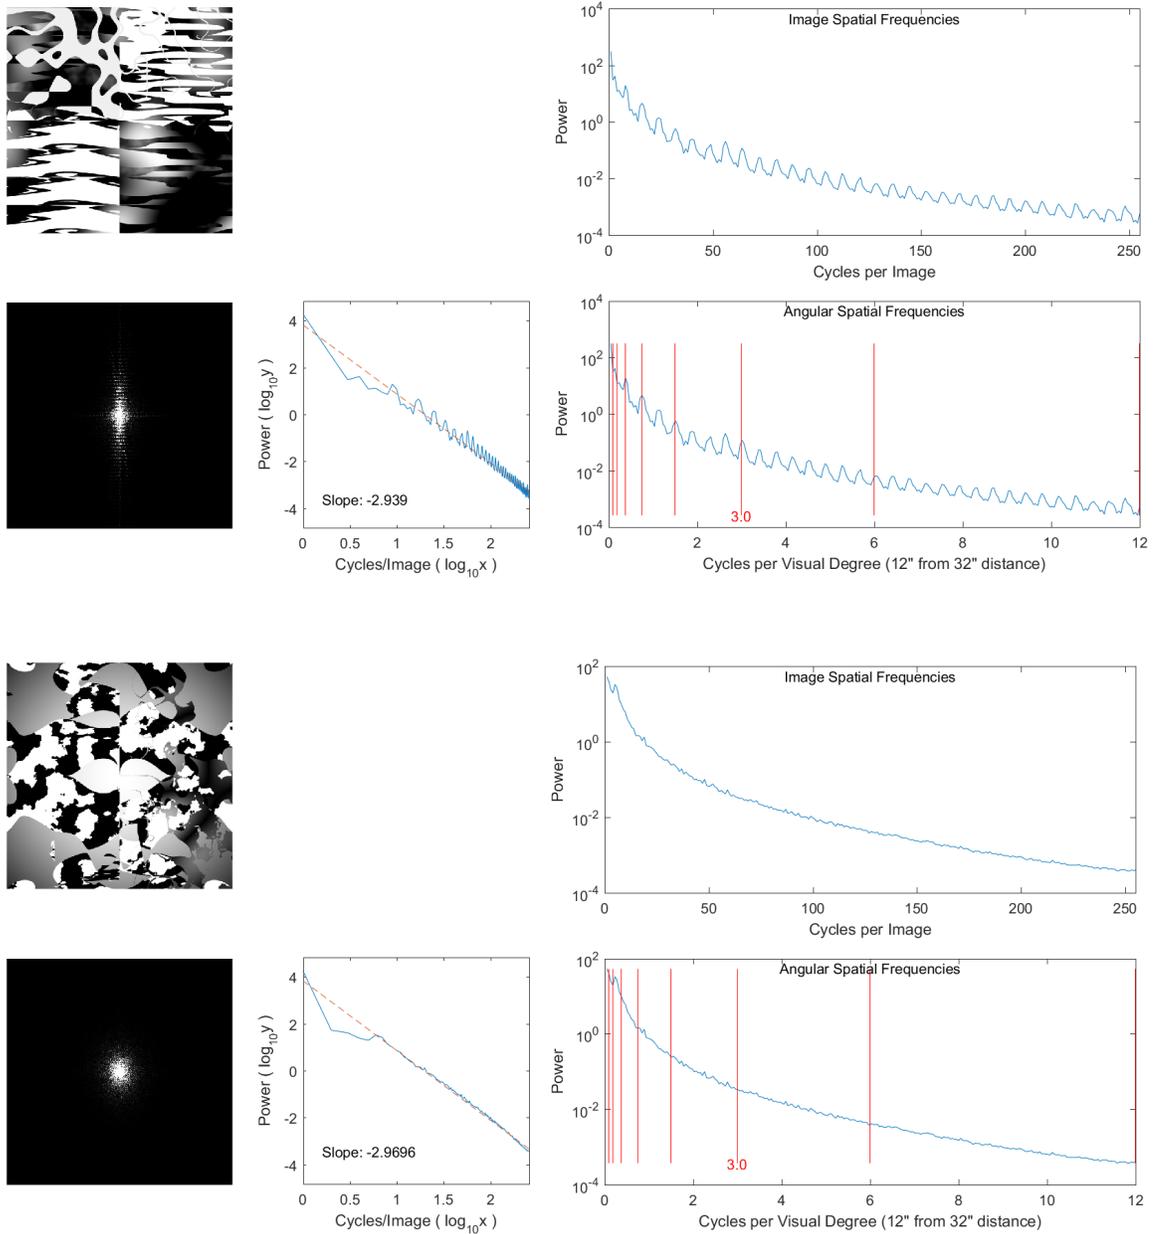
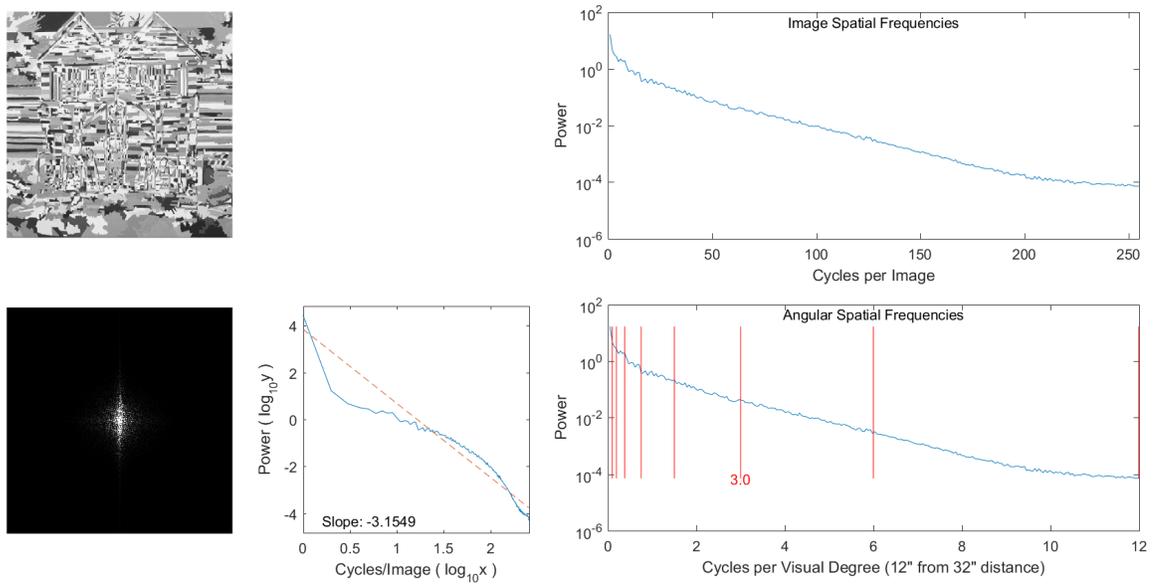
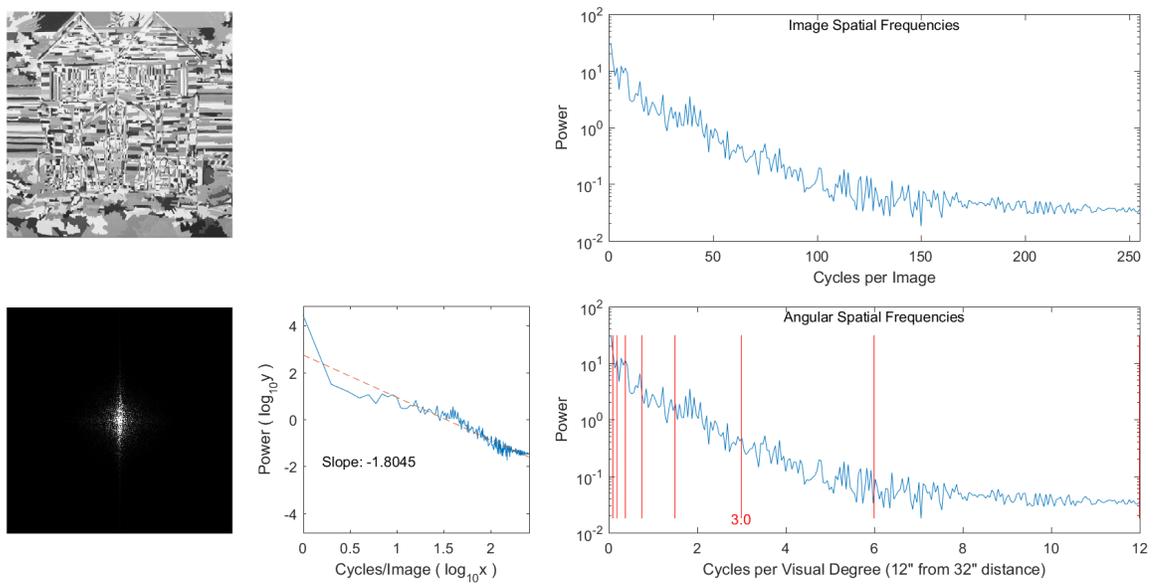


Figure 8.26: Angular Spatial Frequency Analysis: Jesmond Barn

(a) Mapping Across Radially Averaged Power Spectra



(b) Mapping Across Radially Maximized Power Spectra



pleasant effects in the image, adapting the newly created images for purposes outside their considered intent may display unforeseen, less desirable traits. An immediately conceived use case may be with user fitness schemes in evolutionary art systems. With interactive or hybrid fitness depending on user-evaluated thumbnails, large incongruities may appear between the rated thumbnails and full-size renderings.

An additional reservation may be found when considering synthesis of smaller images. As we decrease image size, we will see the 3.0 cycles/degree mark shift toward the lower frequencies within an image. As we have seen that these lower frequencies typically hold higher power, an attempt to reduce power in these frequencies for the goal of aesthetics may provide its own unexplored bias in the search space. We have seen that attempting to change the lower frequency components in a image provides heavy-handed guidance which may greatly change the composition of the images, and negatively affect other objectives if not considered carefully.

The last concern is the most critical, but the one we are now most capable of identifying and accommodating. A 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, but some refinements will be needed otherwise. If provided a 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. We hold high hopes that, with further exploration in the future, novel aesthetic models can be developed from these findings.

# Chapter 9

## Conclusion

### 9.1 Summary

With our background and literature review, compounded with our experimental analysis, we have increased our familiarity with how power spectra relates to the original image spatial attributes.

Our initial, naïve approach of using various power spectra abstractions for guidance was found to be either inadequate, or intractable. While these measures such as the radially averaged power spectra or its regressions may be suitable for providing discrimination during retrieval or classification, they cannot provide sufficient guidance for evolutionary texture synthesis. We find that these approaches allow evolution to quickly and prematurely converge on very basic gradients or shapes.

In attempting to reduce and direct evolutionary pressures, we confirm that the visually salient compositional aspects of an image can be extracted from a small number of power and phase coefficient positions. Specifically, we find that the target-dependent count of coefficients performs well at  $10 \leq K \leq 150$ , with  $K = 50$  performing well for many targets. Optimizations incorporating power spectra symmetry may find good results with  $K$  halved. Experimentation confirms that prioritizing more highly powered frequencies gives a useful method for feature extraction. The truncation and quantization approach suggested by Jacobs et al. [45] shows good utility when adapted for Fourier analysis in an evolutionary application.

Certain targets (such as the Mondrian painting) favoured the more precise, non-quantized error, where others found difficulty in receiving suitable evolutionary guidance unless relaxing error to a count of matching, top-ranked positions. In addition, we found that although certain targets were less dependent on precise phase alignment, most targets could become unrecognisable without consideration of phase.

The precise method of incorporating phase for evolutionary guidance can be target-dependent, and generally difficult.

We recommend the measures employed for experiment P1 (Section 6.3) as a suitable general-purpose compositional guide, which performed well amongst most compositional and genre targets. Suitable guidance to power was found through the truncated, quantized error approach, effectively minimizing the mismatched positions between target and candidate for the highest powered frequencies. The more successful pairing of phase information measured phase error weighted linearly with respect to the position's rank amongst the other top  $K$  coefficients.

Our user survey validates that our general-purpose compositional spatial measure can reproduce images from a target which share visually distinct spatial similarities. We found a statistically significant match between our selected fitness measure and human perception when evaluated amongst the class of basic compositional images. Performance in other target image genres can be improved, and we begin to see increased performance when GP language operators are adjusted for the target.

Where our fitness measures are capable of providing spatial guidance during evolution, language bias and sufficiency cannot be overlooked. Adjustments to GP language and evolution scheme can strongly influence the candidates explored within the evolutionary search space. Minor language adjustments can ease or increase difficulty and tractability for certain spatial traits, image genre, or composition.

When adapting our general-purpose fitness measure for synthesis of colour images, we see an increased difficulty occurring from exacerbated troubles of language sufficiency. Using power and phase measures along the isolated colour channels, we saw success with target images where the intensity maps of each colour channel could have been reproduced as a grayscale target. When incorporating a separate colour distribution objective, such as CHISTQ, we find objective weighting and language considerations are required to prevent shape and colour guidance working in mutual exclusion.

Having identified and analysed a number of less pleasing images produced by our system, our exploration corroborates the findings of Fernandez and Wilkins [41] that certain angular spatial frequencies can be linked to discomfort. Despite a number of concerns having been identified in encoding the phenomenon into an aesthetic model, our exploration with power spectra fitness measures has given us a tool to resolve some of them. We expect that novel aesthetic models may be developed from these findings in the future.

## 9.2 Future Work

In addition to the findings concluded in this report, a large number of avenues remain available for further study. Our guided exploration, while extensive, is by no means exhaustive. A number of alternative approaches and extensions were identified but could not be fully examined, and numerous insights were formed in the analysis of our latter experiments.

The initial exploration with combining both power and phase measures has resulted in phase error being scaled by the significance of its related power coefficient. In developing a multi-objective measure to capture spatial characteristics, the (scaled) weighting between power and phase coefficients has remained largely even. Further refining these weights, particularly in regards to the multi-objective schemes, could benefit from further exploration. The addition of colour measures has shown that phase or power objectives may be sacrificed for colour at a greater scale than expected. Being able to properly tie the phase and power objectives together through scaling or other approaches may be needed to more effectively integrate additional objectives, such as for colour or other aesthetic measures. For possible improvements in producing coloured images, we would propose further study with objective weighting. With the Y+RGB approach in experiment C2, we question if adjustment to weighting for luminance channel objectives could refine the trade-off between shape and colour. Likewise with CHISTQ in experiment C4, it may be interesting to see if heavier CHISTQ weighting and extended generation limits could guide candidates to those with improved colour traits.

Another tool identified to resolve some of the issues regarding Fourier deconstruction with colour channels was the recently proposed Quaternion Fourier Transform (QFT) [34], which has seen use with colour image analysis. While we had hoped to further explore this measure during this research, we had found ourselves unable to give adequate attention to this extended measure, focusing instead on adjustments to the canonical Discrete Fourier Transform. The QFT has shown promise in capturing similar spatial information while overcoming the limitations we observed with colour channel artifacts. We are hopeful that future studies can explore its adaptation for evolving colour image composition, and providing the attention deserved.

An intermediate approach for colour, between independent (RGB) colour channel analysis and the previously mentioned Quaternion Fourier Transform, was the frequency analysis of the image across alternative colour schemes. We had seen certain target perform well when evolved and evaluated using RGB colour channels, but

perform poorly when evaluated using HSL. We suspect that certain targets may see further success when evolved using GP trees mapped to alternative colour channels, such as the mentioned HSL, or  $YC_bC_r$ . Examining these effects with other colour channels, and perhaps identifying optimal colour schemes based on the target image, could be an interesting approach for further research.

While we explored the effectiveness of our measures along a single target for a subset of the possible identified genres, a focussed exploration into each genre could be conducted to evaluate particular effectiveness. We have seen that GP language and windowing can substantially bias how the results are guided through the search space, and the images produced. However, even across visually similar images, the regularity of the image pattern did not appear to substantially affect results (visually or statistically) for simple compositional recreations. A refinement based on these distinctions may show better performance.

Further, each genre may see benefits from domain-specific GP language adjustments and particular refinements away from our general-purpose fitness measures. The language used for our exploration was very simple and algebraic, which leant itself to producing images similar to our compositional targets, but found difficulty in producing the more complex genre targets. The introduction of noise and polar coordinate variables provided a definite bias to the search space, and optimized performance with certain targets. If we are interested in improving similarity with the genre targets, then a more expansive and artistic language may be required. Many of the artistic GP systems explored in the literature offer insights into additional operators and grammars that could be used to this end. We would be most interested in seeing results attempted with more strongly typed, convolutional languages like those shown by Reynolds[72], amongst others.

Prior to settling on our general-purpose fitness measure from experiment P1, we had observed that certain targets showed improved performance with quantized power error measures, where others showed improvement without quantization. Despite the quantized power measures giving comparatively lax guidance in the initial generations (to positive effect), it may be interesting to return to the full and specific error measure after the search space has begun to converge to candidates with approximate shape. Adjusting the fitness measure to this end after a number of generations, or after a certain performance threshold is met, could conceivably provide further refinement to targets which previously performed well using non-quantized power error. We could also find interest in adopting an island model GP approach with populations taking different truncation and quantization schemes for both power and phase.

One factor largely ignored through our exploration was the use of various windows and their subsequent effects on key coefficient extraction. While some trials have been done prior to experimentation, the absence of any windowing on images executed before processing was a decision made to keep in line with existing research found in our literature review. This, however, could still be an oversight persisted from early research, and is itself contradicted from other papers on non-regular signal processing. The effectiveness of our measures could benefit from further study with windowing.

Symmetry presents an existing aesthetic model which has seen successful use in texture synthesis [47]. While not as complete as other methods of measuring symmetry, an analysis of phase angles in powerful coefficient positions can possibly capture mirror symmetries about the origin, as well as tiling aspects. With certain art genres and cultures centred about this aesthetic model [73], power spectra analysis may assist in guiding composition subject to these constraints, and warrants further investigation.

Existing work with wavelet analysis had given large inspiration to a number of our measures' approaches. Abandoning the spatial locality of the wavelet coefficients for the precision of the Fourier amplitudes has provided us a measure of overall compositional similarity, but can be limiting when trying to capture specific subregions of an image. Further comparison between the individual strengths of Fourier and wavelet analysis should be explored in the context of evolutionary texture synthesis. We question if a hybrid model, using measures from both wavelet and Fourier decompositions, could permit evolutions with tunable amounts of compositional and feature replication.

Lastly, we strongly recommend further examination of the uncomfortable image frequencies identified by Fernandez and Wilkins [41]. We have made a cursory examination of the phenomenon in Section 8.3, and formed proposals and hypotheses based on our previous observations with using power spectra measures for evolutionary guidance. However, there are still doubtless many further considerations, refinements, and experiments to be addressed before the phenomenon can be encoded into a useful aesthetic model. We hope to have provided additional resources and encouragement toward its construction, and eagerly await the results of continued research.

# Bibliography

- [1] David S Ebert. *Texturing & modeling: a procedural approach*. Morgan Kaufmann, 2003.
- [2] Marc Hull and Simon Colton. Towards a general framework for program generation in creative domains. In *Proceedings of the 4th International Joint Workshop on Computational Creativity*, pages 137–144, 2007.
- [3] Maryam Baniasadi and Brian J. Ross. Exploring non-photorealistic rendering with genetic programming. *Genetic Programming and Evolvable Machines*, 16(2):211–239, June 2015.
- [4] Craig Reynolds. Interactive evolution of camouflage. *Artificial Life*, 17(2):123–136, Spring 2011.
- [5] Mark Hendrikx, Sebastiaan Meijer, Joeri Van Der Velden, and Alexandru Iosup. Procedural content generation for games: A survey. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 9(1):1–22, 2013.
- [6] John H Holland. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. 1975.
- [7] Karl Sims. Artificial evolution for computer graphics. *ACM Computer Graphics*, 25(4):319–328, July 1991. SIGGRAPH '91 Proceedings.
- [8] D Stork. Computer image analysis of paintings and drawings: An introduction to the literature. In *Proceedings of the Image Processing for Artist Identification Workshop*, 2008.
- [9] Jana Zujovic, Lisa Gandy, Scott Friedman, Bryan Pardo, and Thrasyvoulos N Pappas. Classifying paintings by artistic genre: An analysis of features & clas-

- sifiers. In *Multimedia Signal Processing, 2009. MMSP'09. IEEE International Workshop on*, pages 1–5. IEEE, 2009.
- [10] Dirk Neumann and Karl R Gegenfurtner. Image retrieval and perceptual similarity. *ACM Transactions on Applied Perception (TAP)*, 3(1):31–47, 2006.
- [11] Daniel J Graham and David J Field. Variations in intensity statistics for representational and abstract art, and for art from the eastern and western hemispheres. *Perception*, 37(9):1341–1352, 2008.
- [12] Daniel Graham. Art statistics and visual processing: insights for picture coding. In *Picture Coding Symposium, 2009. PCS 2009*, pages 1–4. IEEE, 2009.
- [13] Daniel J Graham and Christoph Redies. Statistical regularities in art: Relations with visual coding and perception. *Vision Research*, 50(16):1503–1509, 2010.
- [14] Craig Neufeld, Brian J Ross, and William Ralph. The evolution of artistic filters. *The art of artificial evolution*, pages 335–356, 2008.
- [15] Ken Perlin. Improving noise. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 681–682. ACM, 2002.
- [16] Lode Vandevenne. Texture generation using random noise. <http://lodev.org/cgtutor/randomnoise.html>, 2004. [Online; accessed 2017-09-06].
- [17] PM Nishad. Various colour spaces and colour space conversion. *Journal of Global Research in Computer Science (UGC Approved Journal)*, 4(1):44–48, 2013.
- [18] SharkD. Hsl color solid cylinder. [https://commons.wikimedia.org/wiki/File:HSL\\_color\\_solid\\_cylinder\\_alpha\\_lowgamma.png](https://commons.wikimedia.org/wiki/File:HSL_color_solid_cylinder_alpha_lowgamma.png). [Online; accessed 2017-09-18; under Creative Commons Attribution-Share Alike 3.0 Unported <http://creativecommons.org/licenses/by-sa/3.0/deed.en>].
- [19] SharkD. Rgb color solid cube. [https://commons.wikimedia.org/wiki/File:RGB\\_color\\_solid\\_cube.png](https://commons.wikimedia.org/wiki/File:RGB_color_solid_cube.png). [Online; accessed 2017-09-18; under Creative Commons Attribution-Share Alike 3.0 Unported <http://creativecommons.org/licenses/by-sa/3.0/deed.en>].
- [20] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1998.

- [21] Penousal Machado, Juan Romero, Amílcar Cardoso, and Antonino Santos. Partially interactive evolutionary artists. *New Generation Computing*, 23(2):143–155, 2005.
- [22] Simon Colton, Michael Cook, and Azalea Raad. Ludic considerations of tablet-based evo-art. In *Applications of Evolutionary Computation*, pages 223–233. Springer, 2011.
- [23] Jinhong Zhang, Rasmus Taarnby, Antonios Liapis, and Sebastian Risi. Draw-compileevolve: Sparking interactive evolutionary art with human creations. In *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, pages 261–273. Springer, 2015.
- [24] D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, 1989.
- [25] 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 Verlag, 1997.
- [26] D. Corne and J. Knowles. Techniques for highly multiobjective optimisation: Some nondominated points are better than others. In *Proc. GECCO 2007*, pages 773–780. ACM Press, 2007.
- [27] Brian J. Ross and Han Zhu. Procedural texture evolution using multiobjective optimization. *New Generation Computing*, 22(3):271–293, 2004.
- [28] Brian J. Ross, William Ralph, and Hai Zong. Evolutionary image synthesis using a model of aesthetics. In Gary G. Yen, Lipo Wang, Piero Bonissone, and Simon M. Lucas, editors, *Proceedings of the 2006 IEEE Congress on Evolutionary Computation*, pages 3832–3839, Vancouver, 6-21 July 2006. IEEE Press.
- [29] Pierre Brémaud. *Mathematical principles of signal processing: Fourier and wavelet analysis*. Springer Science & Business Media, 2013.
- [30] Ren Schwarz. Fourier synthesis. [https://commons.wikimedia.org/wiki/File:Fourier\\_synthesis.svg](https://commons.wikimedia.org/wiki/File:Fourier_synthesis.svg). [Online; accessed 2017-09-18; under Creative Commons Attribution-Share Alike 3.0 Unported <http://creativecommons.org/licenses/by-sa/3.0/deed.en>].

- [31] Alan Gleason. *Who is Fourier?: a mathematical adventure*. Language Research Foundation, 1995.
- [32] André Rauh and Gonzalo R Arce. Sparse 2d fast fourier transform. In *Proc. 10th Int. Conf. Sampling Theory Appl.*, 2012.
- [33] Stephen J Sangwine. The problem of defining the fourier transform of a colour image. In *1998. ICIP 98. Proceedings. International Conference on Image Processing*, volume 1, pages 171–175. IEEE, 1998.
- [34] Vikas R Dubey. Quaternion fourier transform for colour images. *International Journal Computer Science and Information Technology*, 5(3), 2014.
- [35] National Semiconductor. Power spectra estimation. <http://www.dcs.warwick.ac.uk/~feng/teaching/PowerSpectrum.pdf>. [Online; accessed 2018-03-01].
- [36] John M. Brayer. Introductions to fourier transforms for image processing. <https://www.cs.unm.edu/~brayer/vision/fourier.html>. [Online; accessed 2017-09-15].
- [37] Amara Graps. An introduction to wavelets. *IEEE Computational Science and Engineering*, 2(2):50–61, 1995.
- [38] Murray B Sachs, Jacob Nachmias, and John G Robson. Spatial-frequency channels in human vision. *JOSA*, 61(9):1176–1186, 1971.
- [39] Lamberto Maffei and Adriana Fiorentini. The visual cortex as a spatial frequency analyser. *Vision research*, 13(7):1255–1267, 1973.
- [40] Patrik Vuilleumier, Jorge L Armony, Jon Driver, and Raymond J Dolan. Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature neuroscience*, 6(6):624, 2003.
- [41] Dominic Fernandez and Arnold J Wilkins. Uncomfortable images in art and nature. *Perception*, 37(7):1098–1113, 2008.
- [42] Rosario M Balboa and Norberto M Grzywacz. Power spectra and distribution of contrasts of natural images from different habitats. *Vision Research*, 43(24):2527–2537, 2003.

- [43] Alaa Eldin M Ibrahim. *Genshade: an evolutionary approach to automatic and interactive procedural texture generation*. PhD thesis, Texas A&M University, 1998.
- [44] Andrea L. Wiens and Brian J. Ross. Gentropy: Evolutionary 2D texture generation. In Darrell Whitley, editor, *Late Breaking Papers at the 2000 Genetic and Evolutionary Computation Conference*, pages 418–424, Las Vegas, Nevada, USA, 8 July 2000.
- [45] Charles E Jacobs, Adam Finkelstein, and David H Salesin. Fast multiresolution image querying. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, pages 277–286. ACM, 1995.
- [46] Shumeet Baluja, Dean Pomerleau, and Todd Jochem. Towards automated artificial evolution for computer-generated images. *Connection Science*, 6(2 and 3):325–354, 1994.
- [47] Eelco den Heijer and A. E. Eiben. Investigating aesthetic measures for unsupervised evolutionary art. *Swarm and Evolutionary Computation*, 16:52–68, 2014.
- [48] Yang Li, Changjun Hu, Leandro L. Minku, and Haolei Zuo. Learning aesthetic judgements in evolutionary art systems. *Genetic Programming and Evolvable Machines*, 14(3):315–337, September 2013. Special issue on biologically inspired music, sound, art and design.
- [49] Roisin Loughran, James McDermott, and Michael O’Neill. Tonality driven piano compositions with grammatical evolution. In Yadahiko Murata, editor, *Proceedings of 2015 IEEE Congress on Evolutionary Computation (CEC 2015)*, pages 2168–2175, Sendai, Japan, 25-28 May 2015. IEEE Press.
- [50] Steve Bergen and Brian J. Ross. Aesthetic 3D model evolution. *Genetic Programming and Evolvable Machines*, 14(3):339–367, September 2013. Special issue on biologically inspired music, sound, art and design.
- [51] Long Nguyen, Daniel Lang, Nico van Gessel, Anna K Beike, Achim Menges, Ralf Reski, and Anita Roth-Nebelsick. Evolutionary processes as models for exploratory design. In *Biomimetic Research for Architecture and Building Construction*, pages 295–318. Springer, 2016.

- [52] Rafal Kicinger, Tomasz Arciszewski, and Kenneth De Jong. Evolutionary computation and structural design: A survey of the state-of-the-art. *Computers & Structures*, 83(23):1943–1978, 2005.
- [53] Adriana Fiorentini and Nicoletta Berardi. Perceptual learning specific for orientation and spatial frequency. *Nature*, 287(5777):43–44, 1980.
- [54] Russell L De Valois, Duane G Albrecht, and Lisa G Thorell. Spatial frequency selectivity of cells in macaque visual cortex. *Vision research*, 22(5):545–559, 1982.
- [55] RP Millane, S Alzaidi, and WH Hsiao. Scaling and power spectra of natural images. In *Proc. Image and Vision Computing New Zealand*, pages 148–153, 2003.
- [56] Daniel J Graham, Jay D Friedenber, and Daniel N Rockmore. Efficient visual system processing of spatial and luminance statistics in representational and non-representational art. In *IS&T/SPIE 7240 Electronic Imaging XIV*. International Society for Optics and Photonics, 2009.
- [57] Carlton W Niblack, Ron Barber, Will Equitz, Myron D Flickner, Eduardo H Glasman, Dragutin Petkovic, Peter Yanker, Christos Faloutsos, and Gabriel Taubin. Qbic project: querying images by content, using color, texture, and shape. In *IS&T/SPIE's Symposium on Electronic Imaging: Science and Technology*, pages 173–187. International Society for Optics and Photonics, 1993.
- [58] Bela Julesz and Terry Caelli. On the limits of fourier decompositions in visual texture perception. *Perception*, 8(1):69–73, 1979.
- [59] Inc. The MathWorks. Mathworks: Matlab. <https://www.mathworks.com/>.
- [60] Sean Luke, Liviu Panait, Gabriel Balan, Sean Paus, Zbigniew Skolicki, Jeff Bassett, Robert Hubley, and A Chircop. Ecj: A java-based evolutionary computation research system. *Downloadable versions and documentation can be found at the following url: <http://cs.gmu.edu/eclab/projects/ecj>*, 2006.
- [61] Steven Bergen and Brian J. Ross. Evolutionary art using summed multi-objective ranks. In Rick Riolo, Trent McConaghy, and Ekaterina Vladislavleva, editors, *Genetic Programming Theory and Practice VIII*, volume 8 of *Genetic and Evolutionary Computation*, chapter 14, pages 227–244. Springer, Ann Arbor, USA, 20-22 May 2010.

- [62] Riven. “perlinnoise :: smooth/turbulent”. <http://riven8192.blogspot.ca/2009/08/perlinnoise.html>. [Online; accessed 2017-08-03; under Creative Commons Attribution 3.0 Unported <http://creativecommons.org/licenses/by/3.0/>].
- [63] Stefan Gustavson. Simplex noise demystified. <http://staffwww.itn.liu.se/~stegu/simplexnoise/simplexnoise.pdf>. [Online; accessed 2017-08-03].
- [64] Procsilas Moscas. Cable ends - san francisco bay bridge. <https://www.flickr.com/photos/procsilas/11315953/>. [Online; accessed 2017-08-03; under Creative Commons Attribution 2.0 Generic <http://creativecommons.org/licenses/by/2.0/>].
- [65] Paul Walsh and Prasad Gade. The use of an aesthetic measure for the evolution of fractal landscapes. In *Evolutionary Computation (CEC), 2011 IEEE Congress on*, pages 1613–1619. IEEE, 2011.
- [66] Elham Salimi. Statistical image analysis for image evolution. Master’s thesis, Brock University, 2016.
- [67] Dennis Wackerly, William Mendenhall, and Richard Scheaffer. *Mathematical statistics with applications*. Nelson Education, 2007.
- [68] Jean Dickinson Gibbons and Subhabrata Chakraborti. *Nonparametric statistical inference*. Springer, 2011.
- [69] Christophe Dang Ngoc Chan. Coordonnes circulaires (coordonnes polaires dans le plan). [https://commons.wikimedia.org/wiki/File:Coordonnees\\_polaires\\_plan.png](https://commons.wikimedia.org/wiki/File:Coordonnees_polaires_plan.png). [Online; accessed 2017-09-18; under Creative Commons Attribution-Share Alike 3.0 Unported <http://creativecommons.org/licenses/by-sa/3.0/deed.en>].
- [70] Steven Bergen and Brian J Ross. Automatic and interactive evolution of vector graphics images with genetic algorithms. *The Visual Computer*, 28(1):35–45, 2012.
- [71] Naoum P Issa, Christopher Trepel, and Michael P Stryker. Spatial frequency maps in cat visual cortex. *Journal of Neuroscience*, 20(22):8504–8514, 2000.
- [72] Craig Reynolds. Evolving textures from high level descriptions: Gray with an accent color. In *Applications of Evolutionary Computation*, pages 384–393. Springer, 2011.

- [73] Ameneh McCullough. Symmetry and islamic art. 2014.

# Appendix A

## Target Analysis

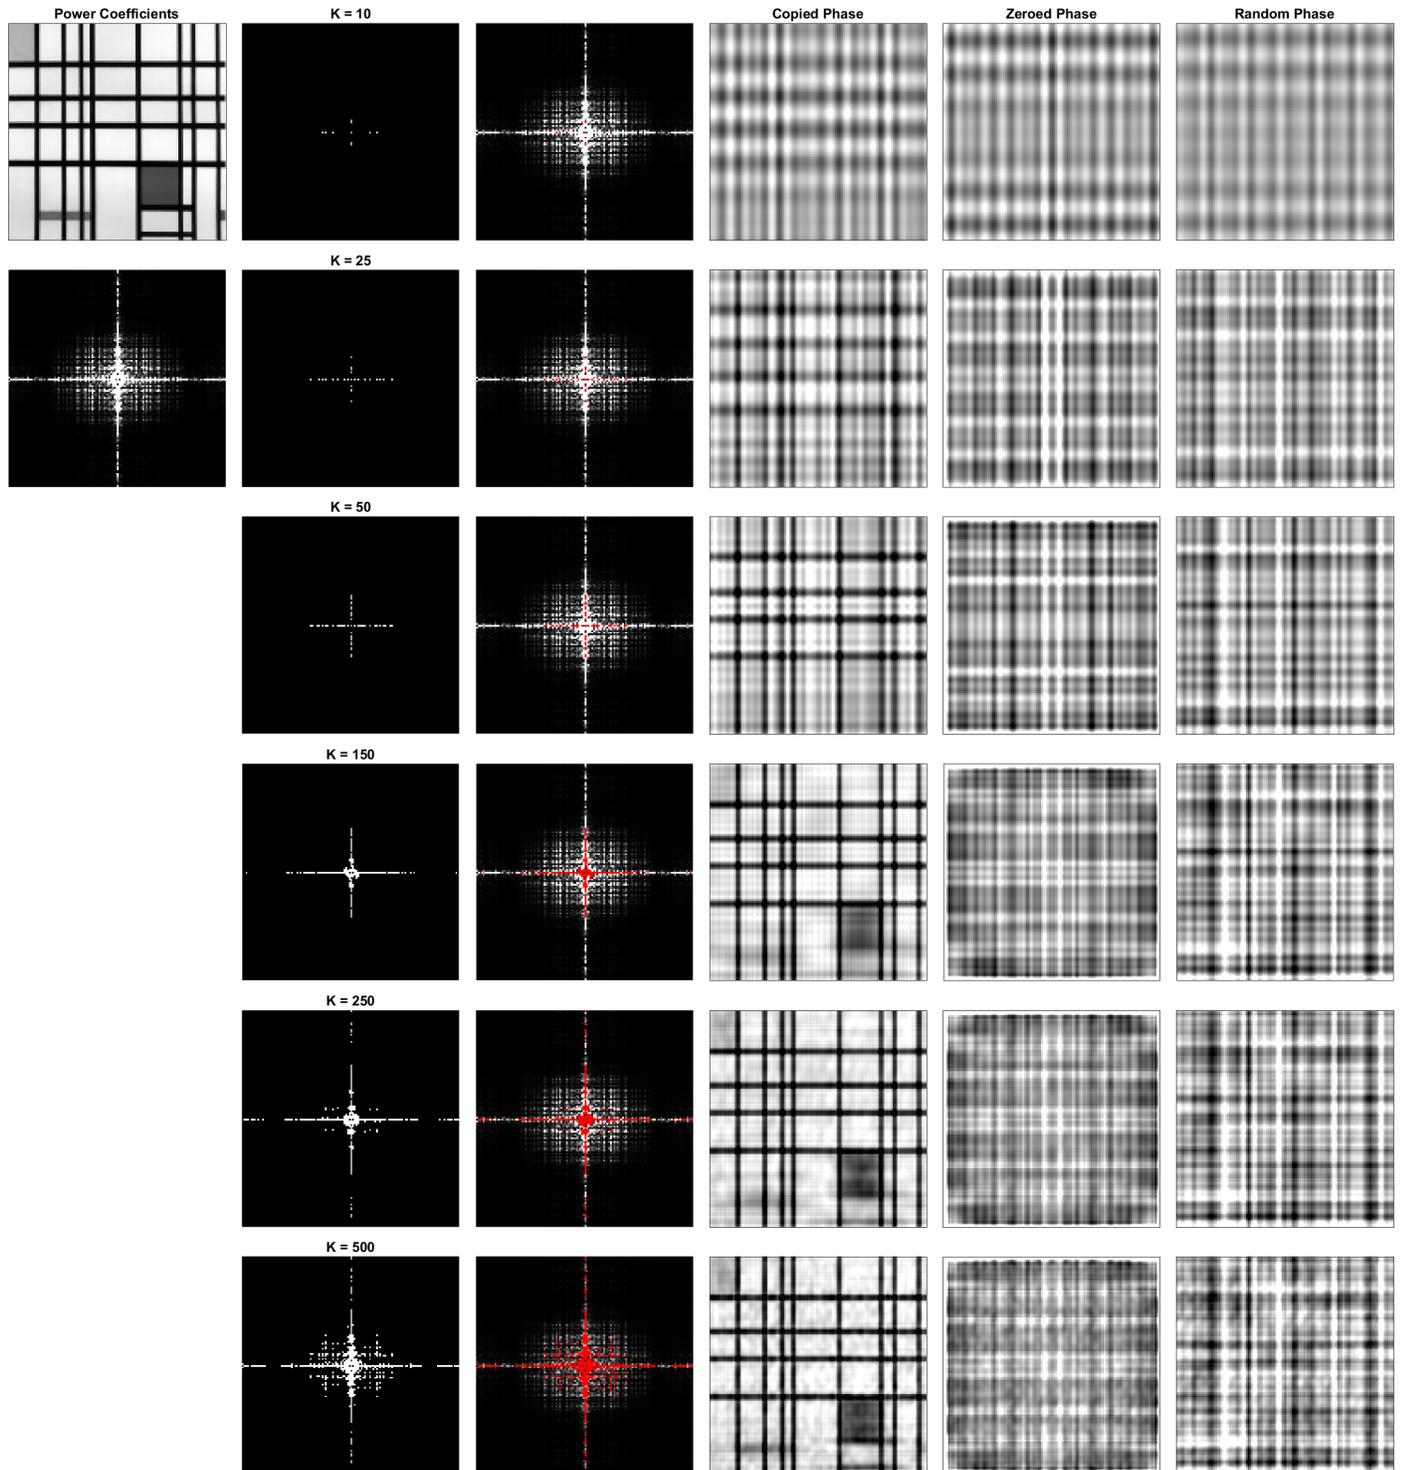
Below, in each of the first (leftmost) columns, we have the various targets used in our explorations showed with their respective 2D Fourier decomposition power coefficients. We then consider isolating the top  $K$  coefficients of the target, and reconstructing the image using only those coefficients.

The second column shows the locations of the isolated coefficients. The third column shows the isolated coefficients in red and overlaid to better show their placements in the full power decomposition grid. The fourth column uses the inverse Fourier transform to reconstruct an image using only the magnitude and phase of the  $K$  isolated coefficient positions. The fifth and sixth columns reconstruct an image similar to the fourth column, but using a constant zero phase value, or random phase values respectively.

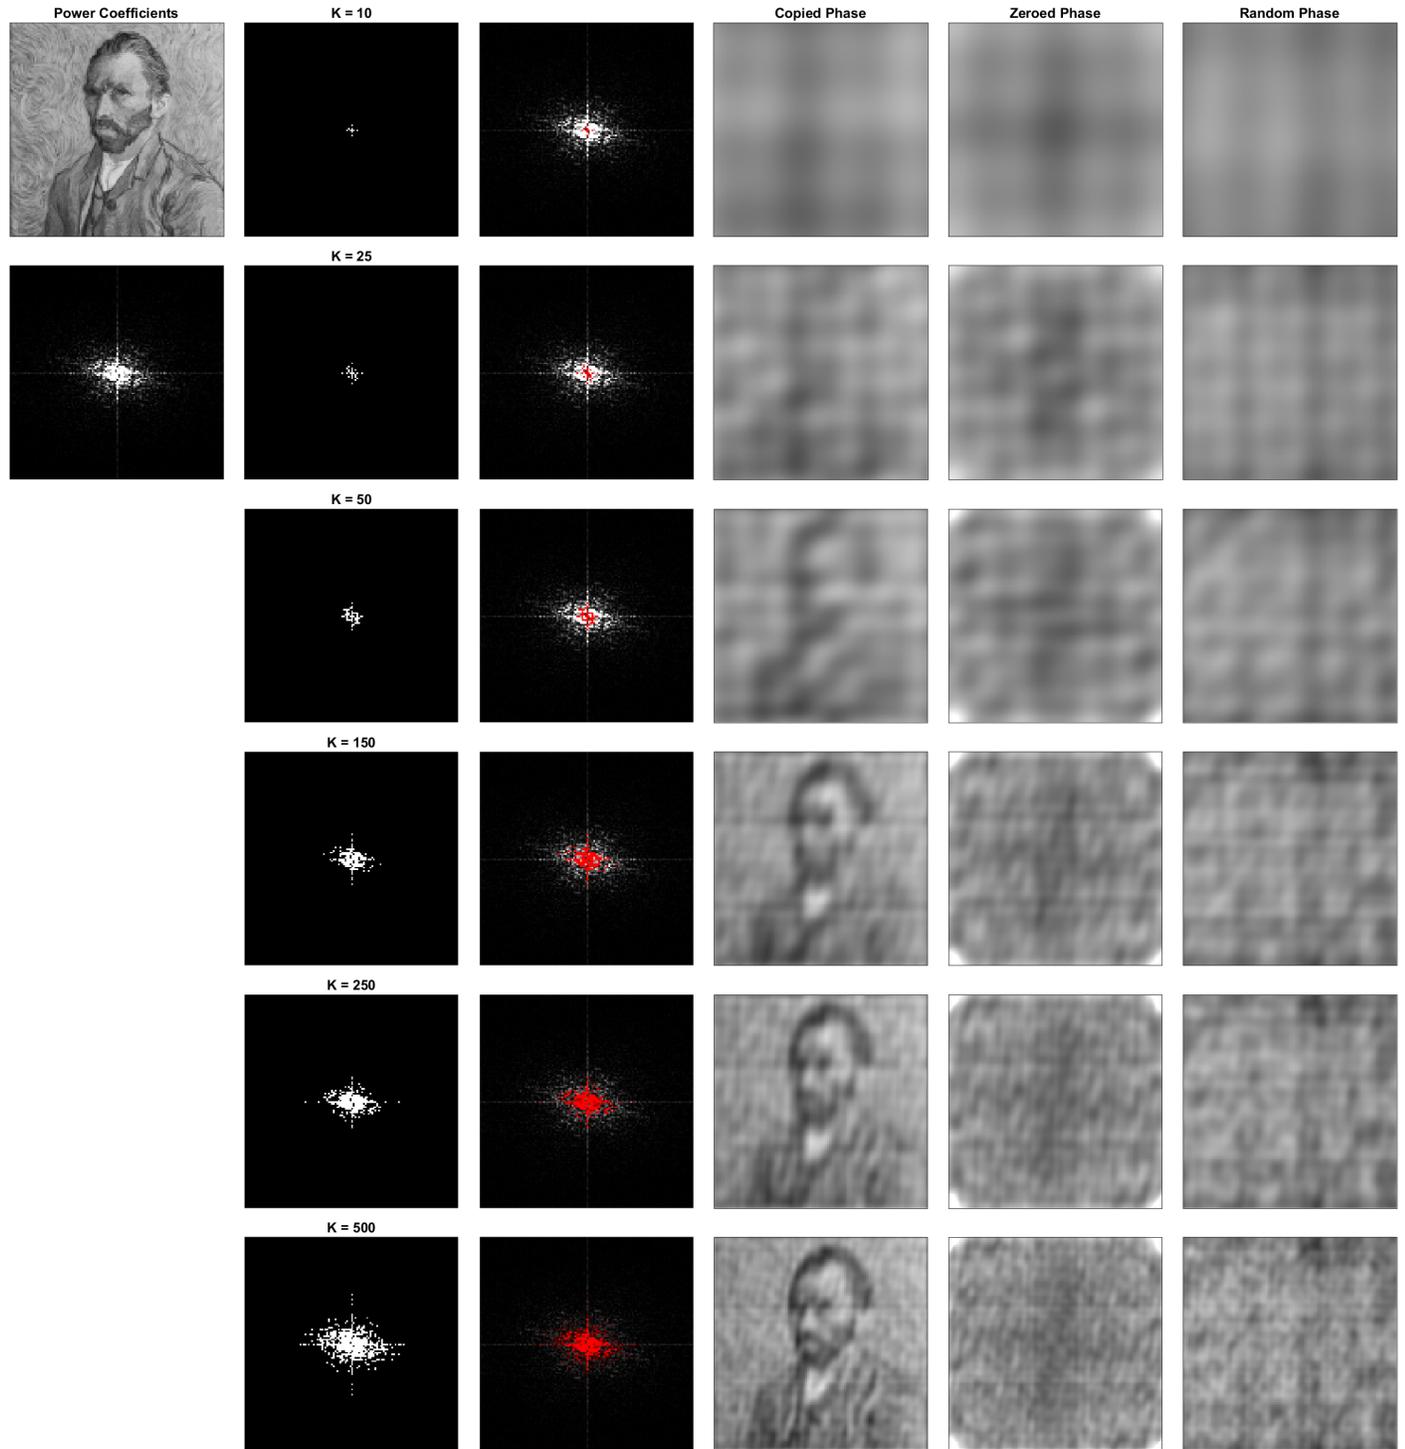
The randomized phase values persist across all values of  $K$  for a given target image in the below illustrations.

Figure A.1: Target image analysis with FFT decompositions and reconstructions

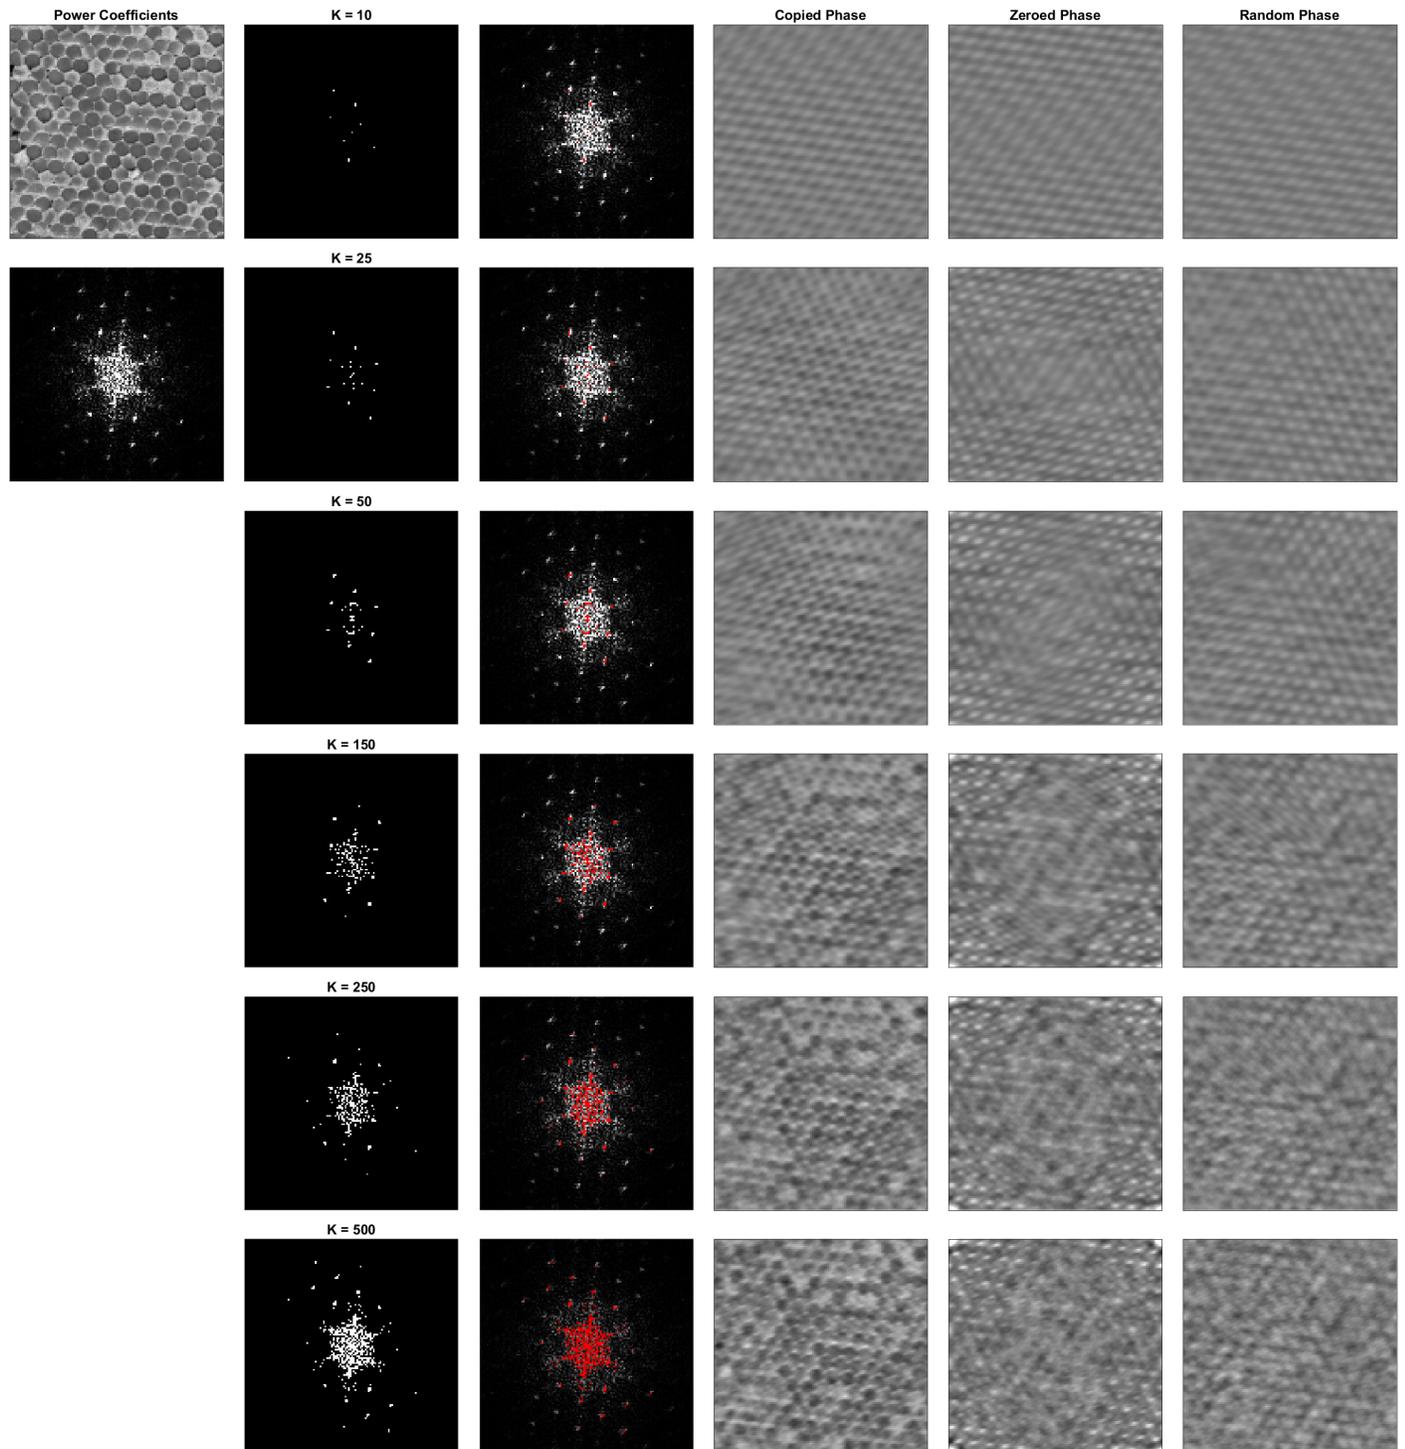
(a) Abstract Painting - Mondrian



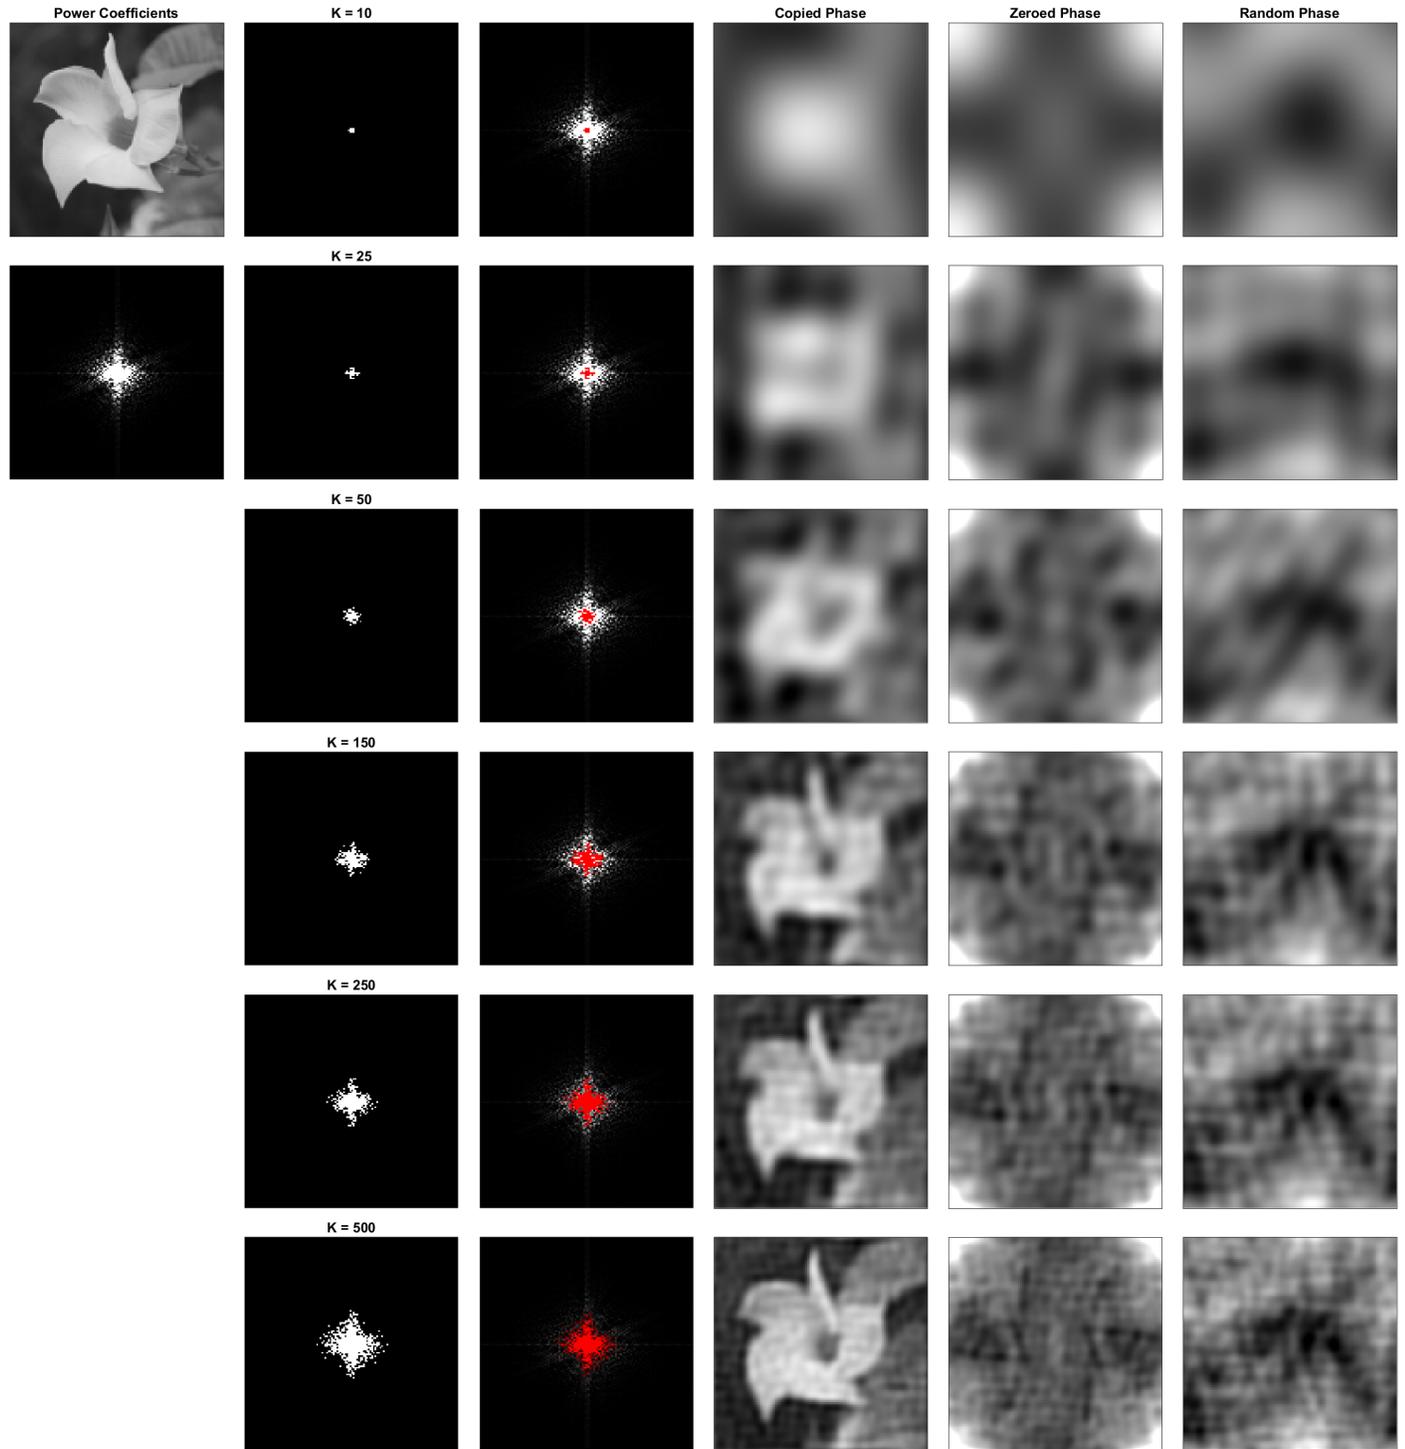
(b) Portrait Painting - Van Gogh



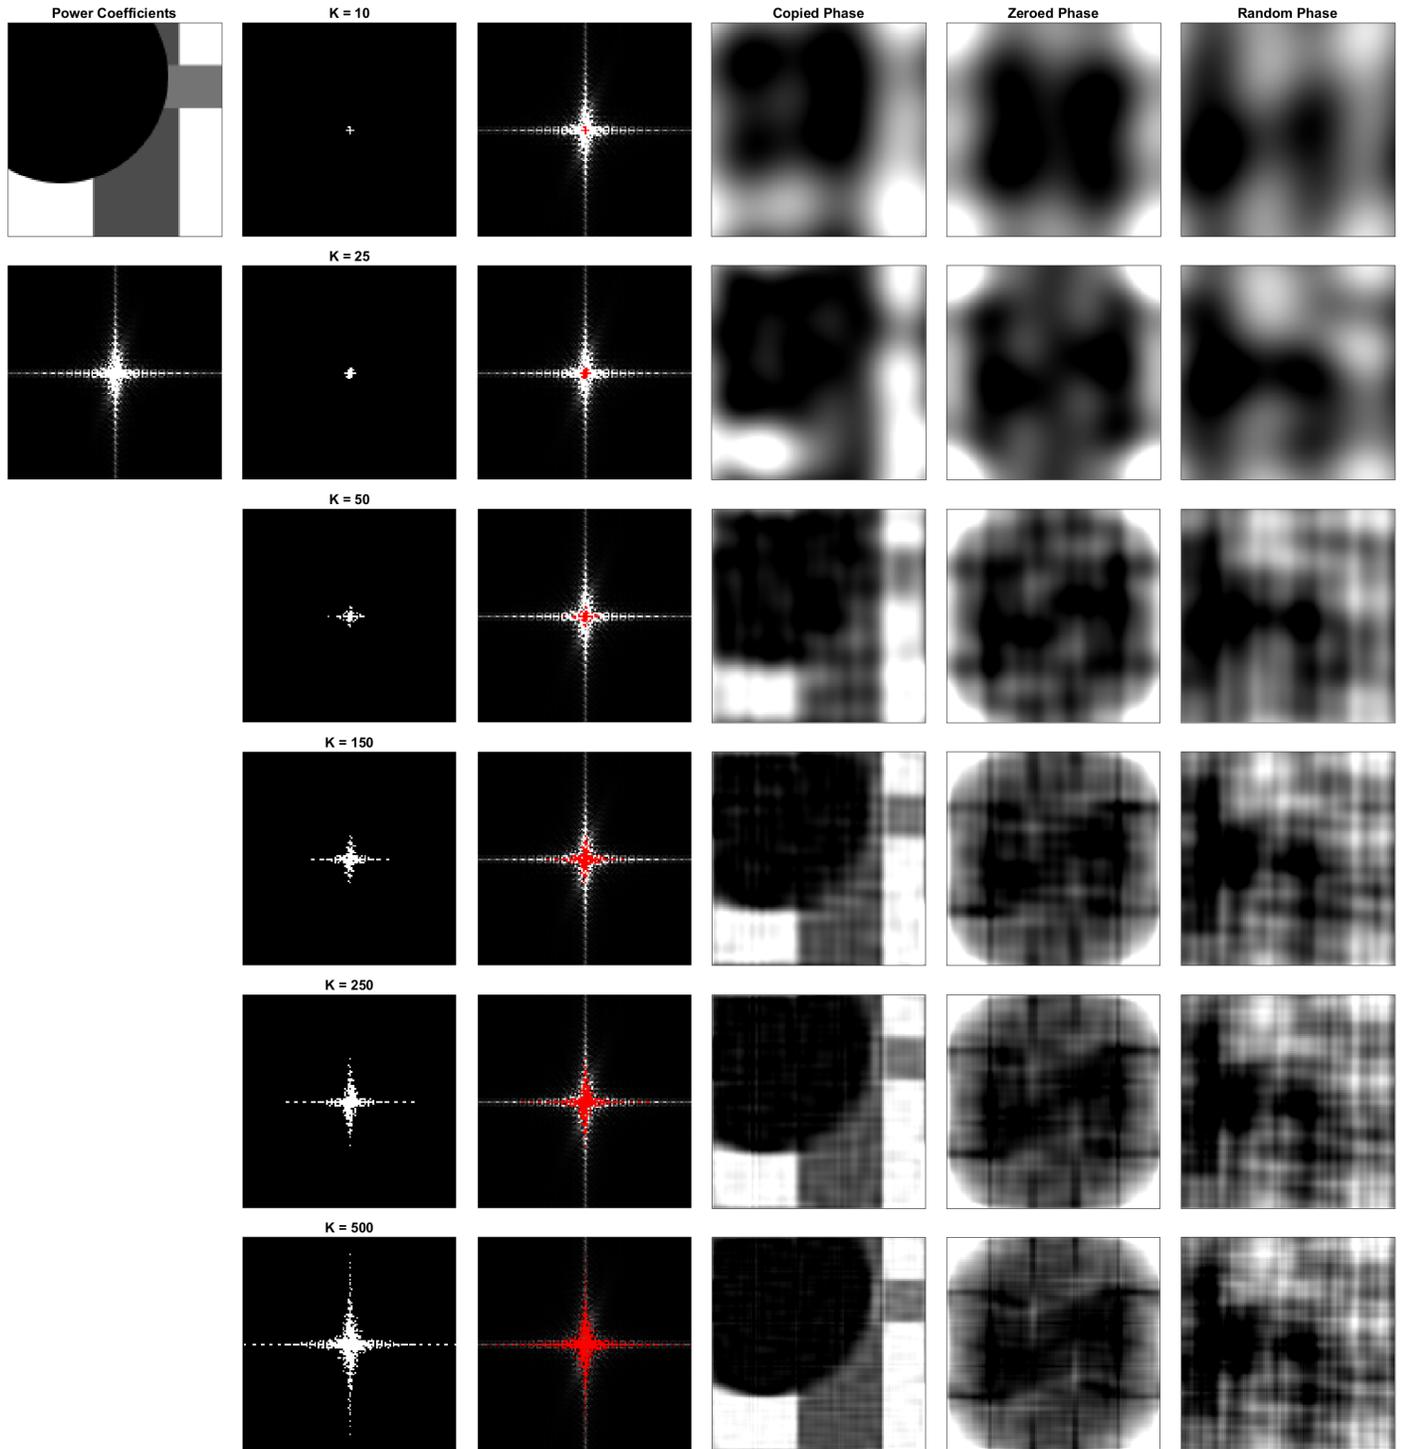
(c) Texture Photo - Cable Cross-Section



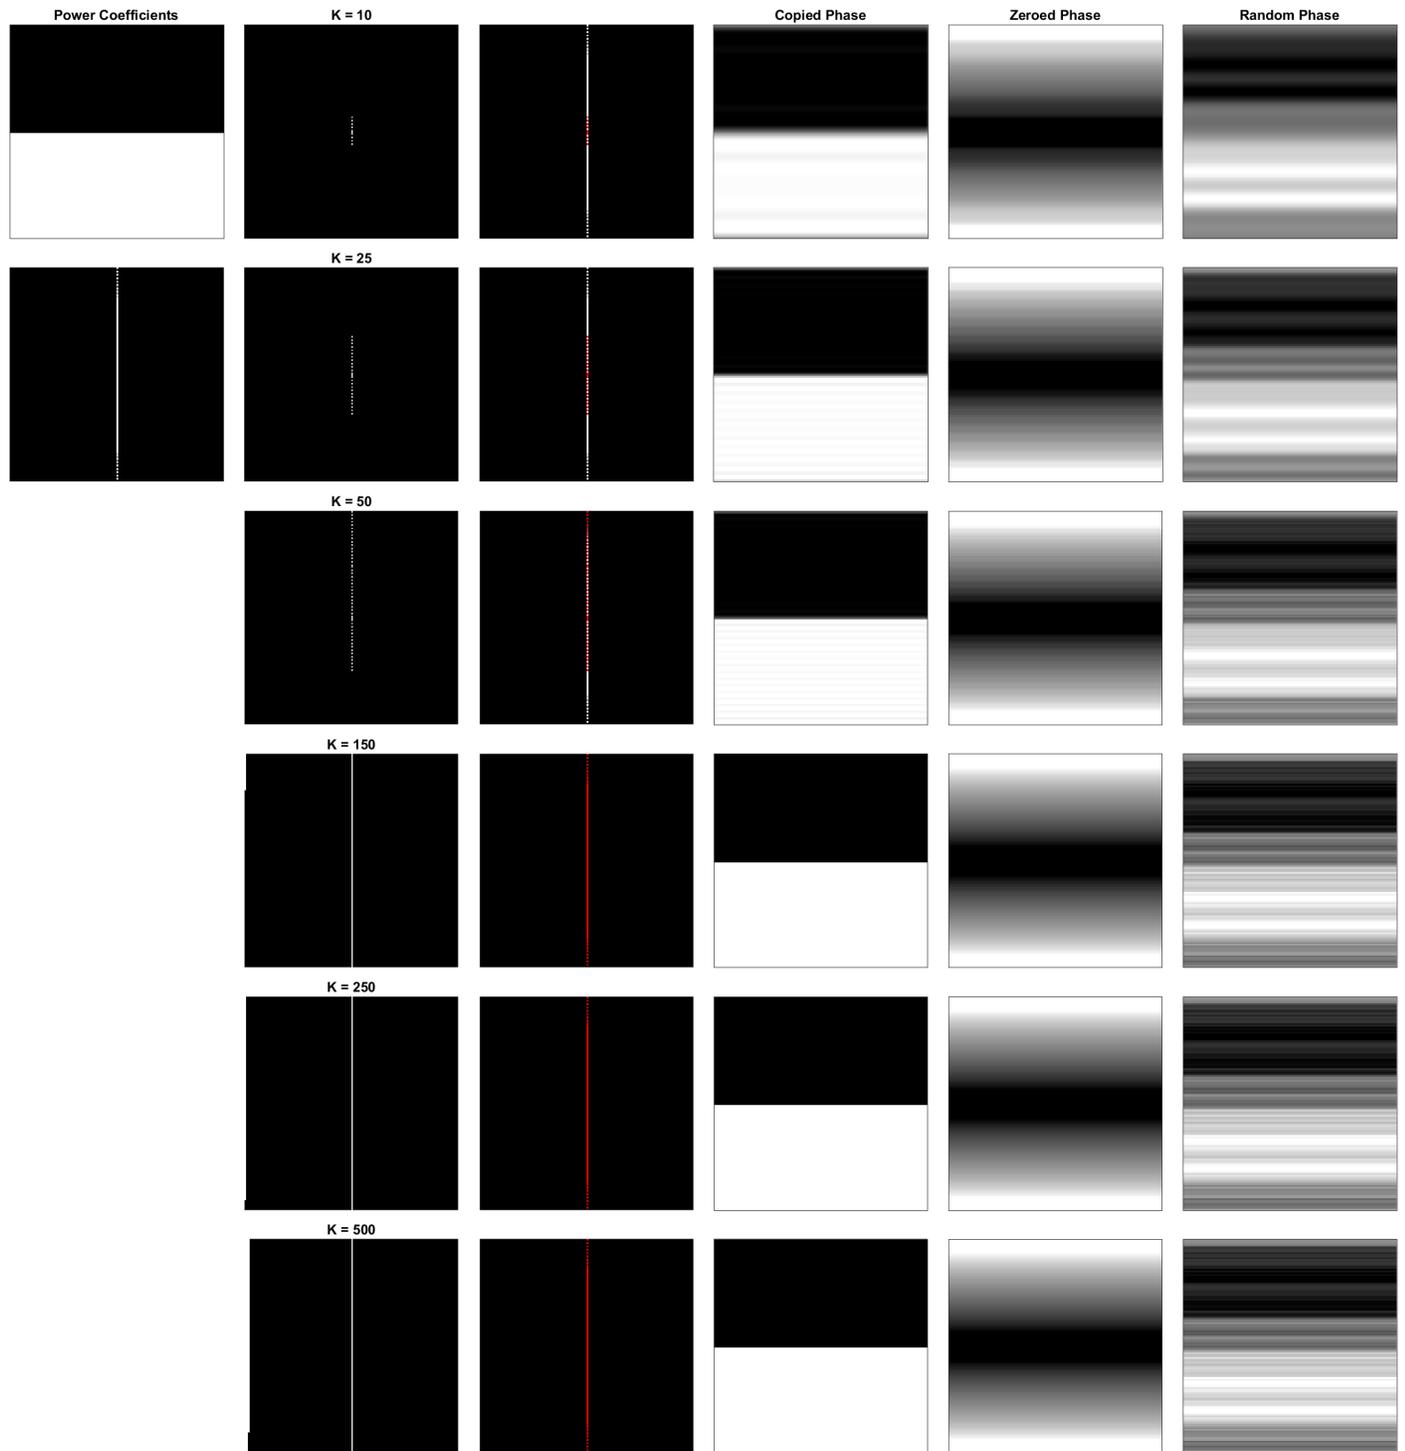
(d) Object Photo - Flower



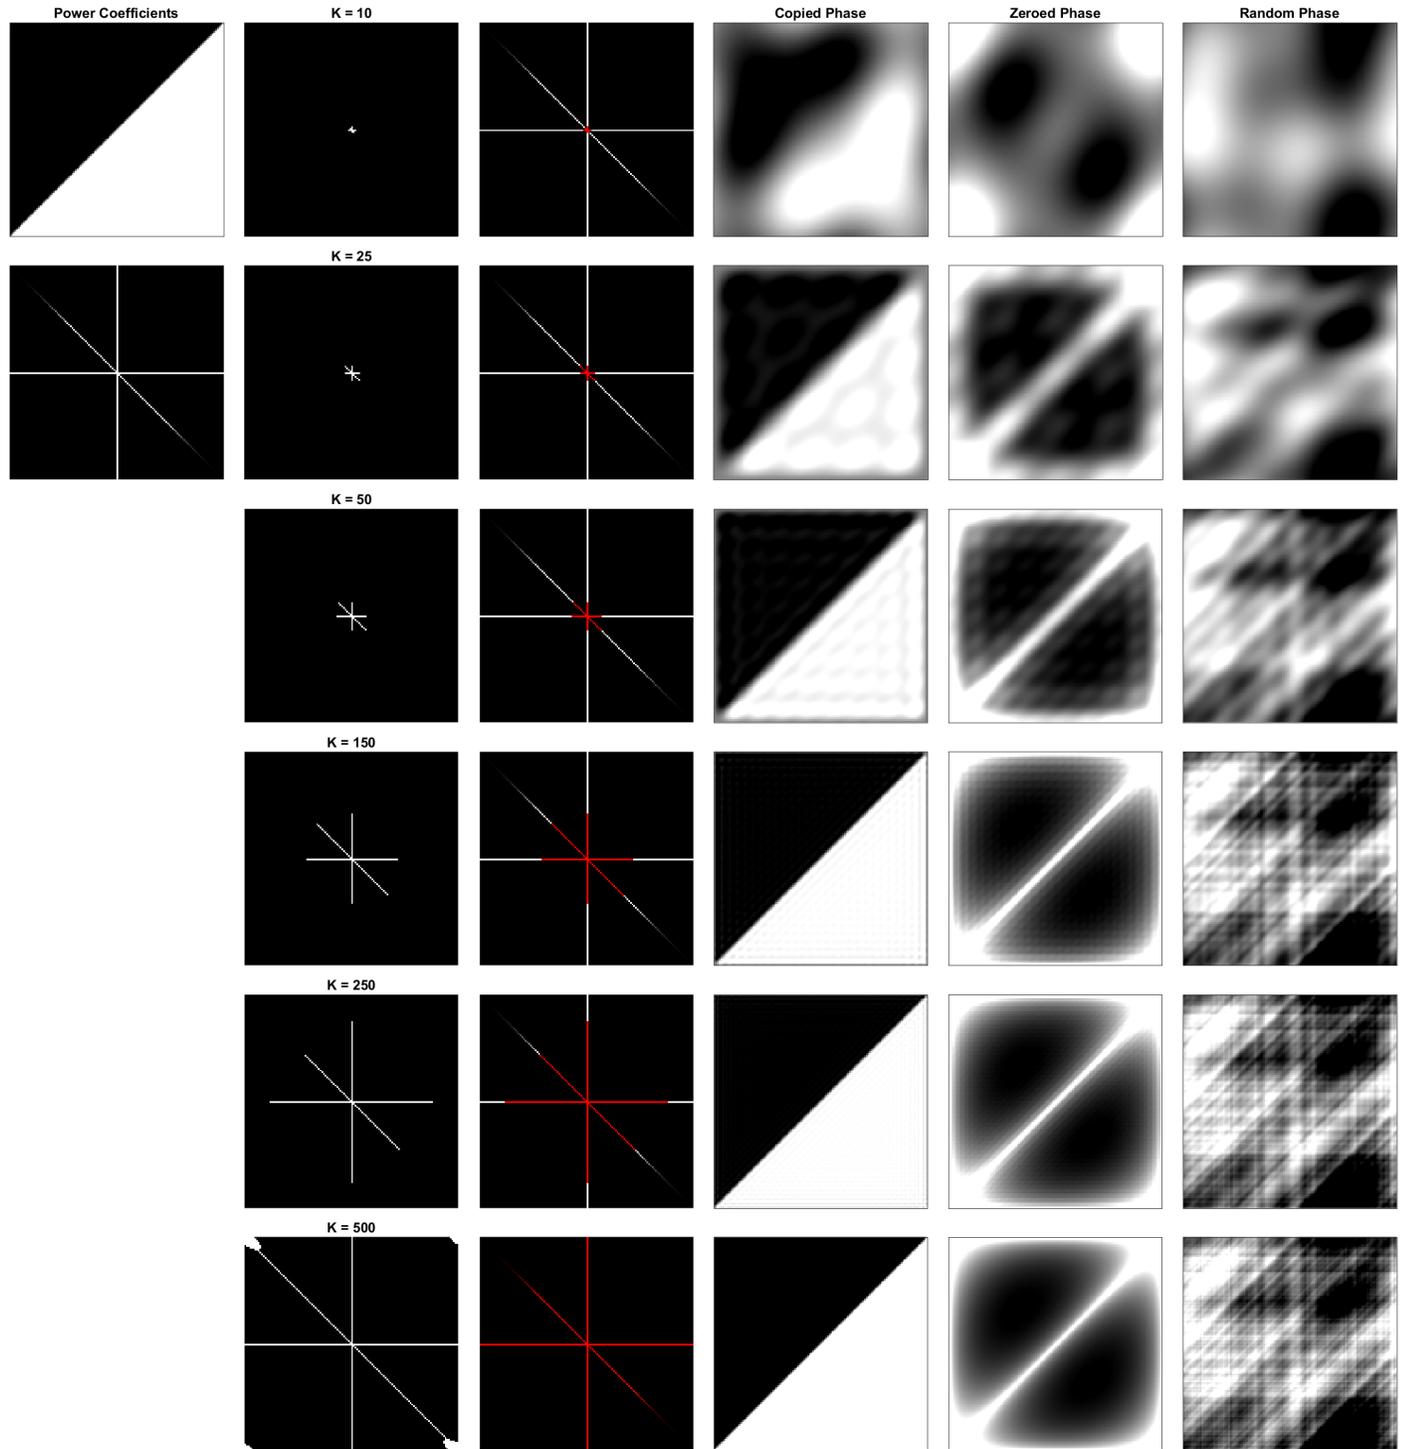
(e) Composition 01



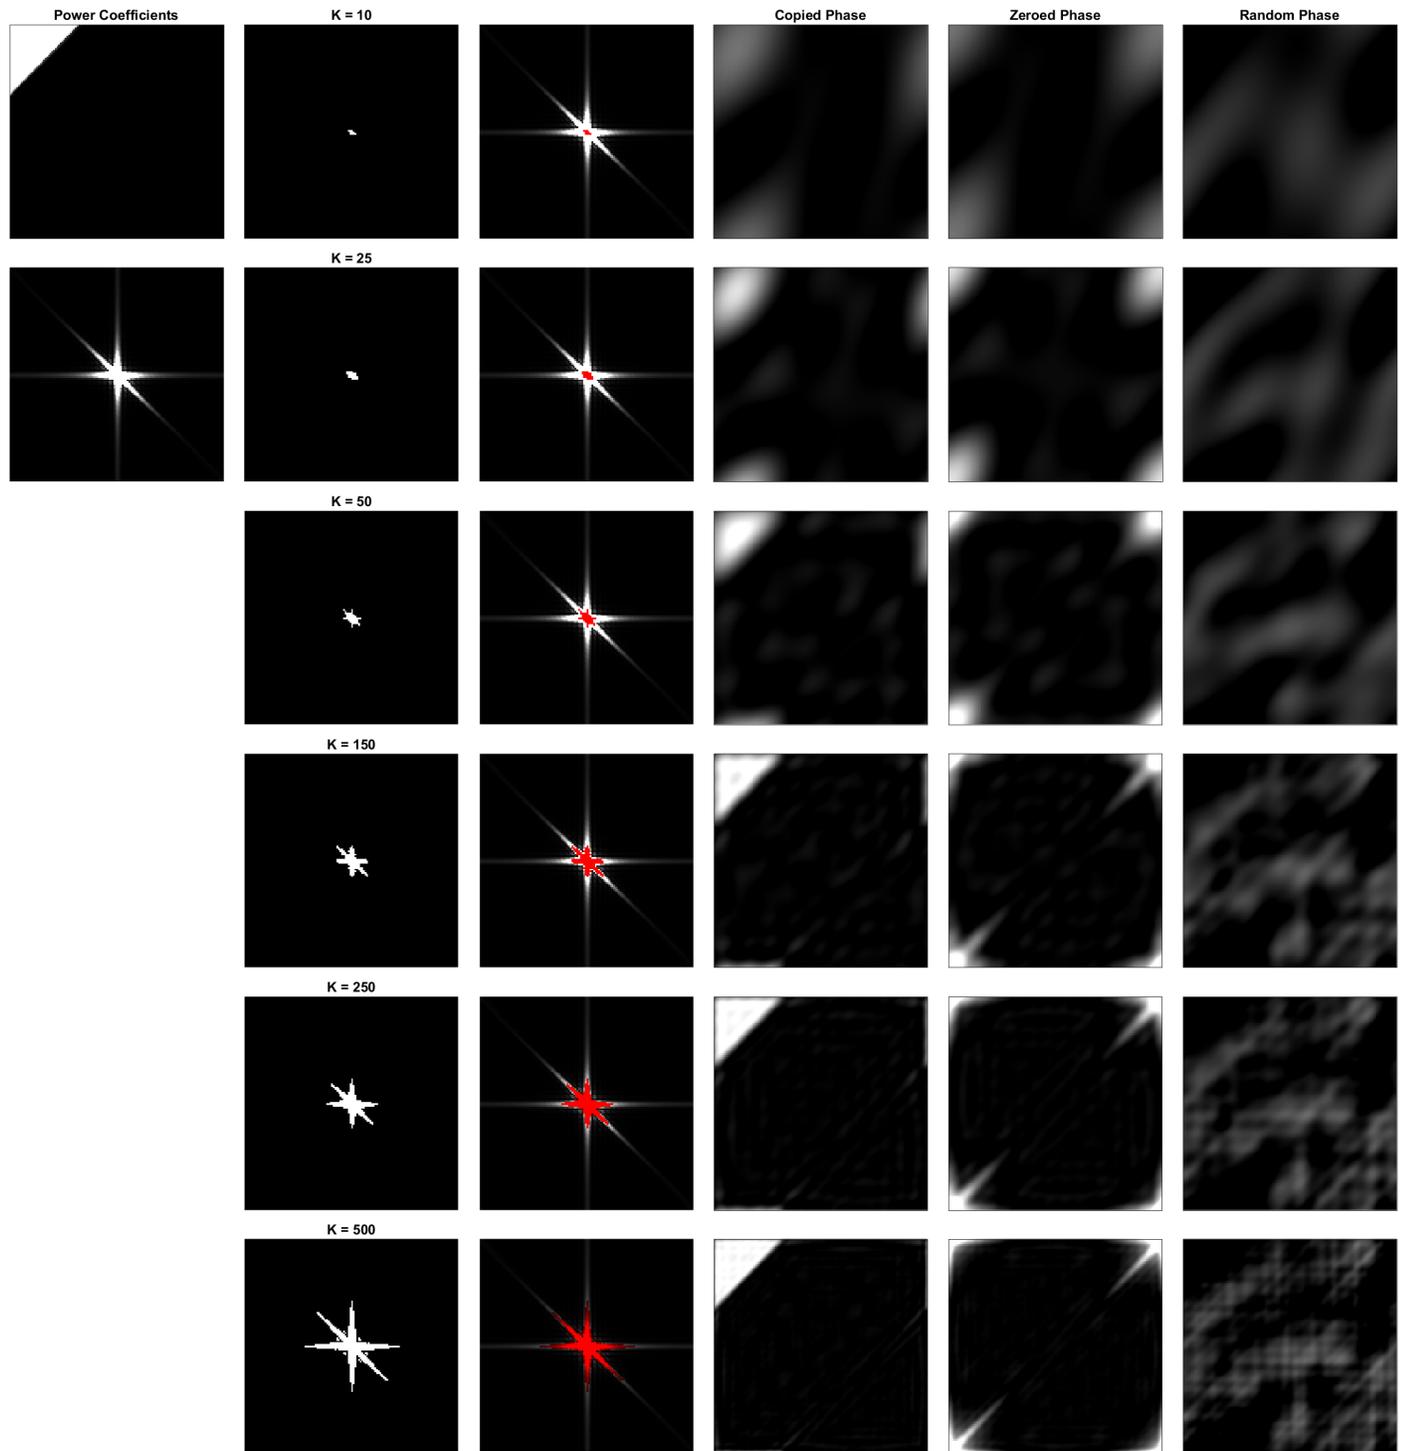
(f) Composition 02



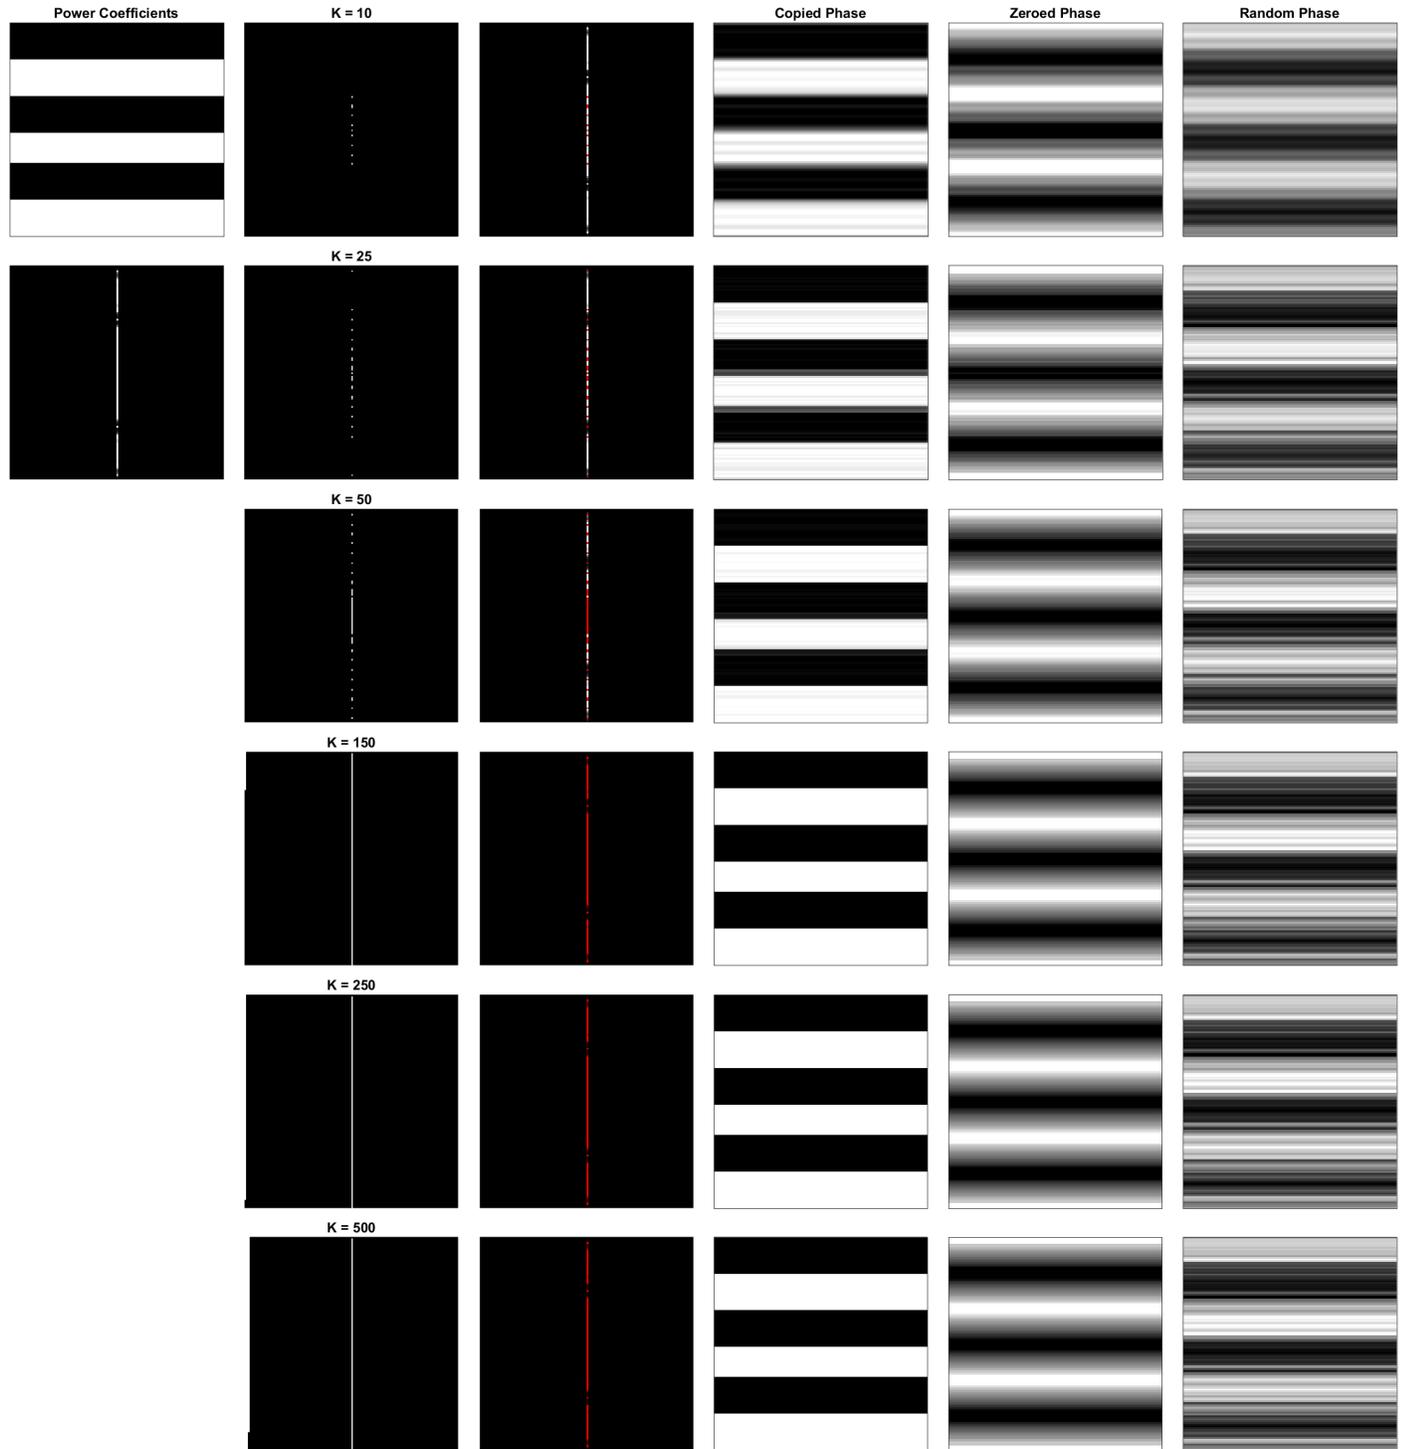
(g) Composition 03



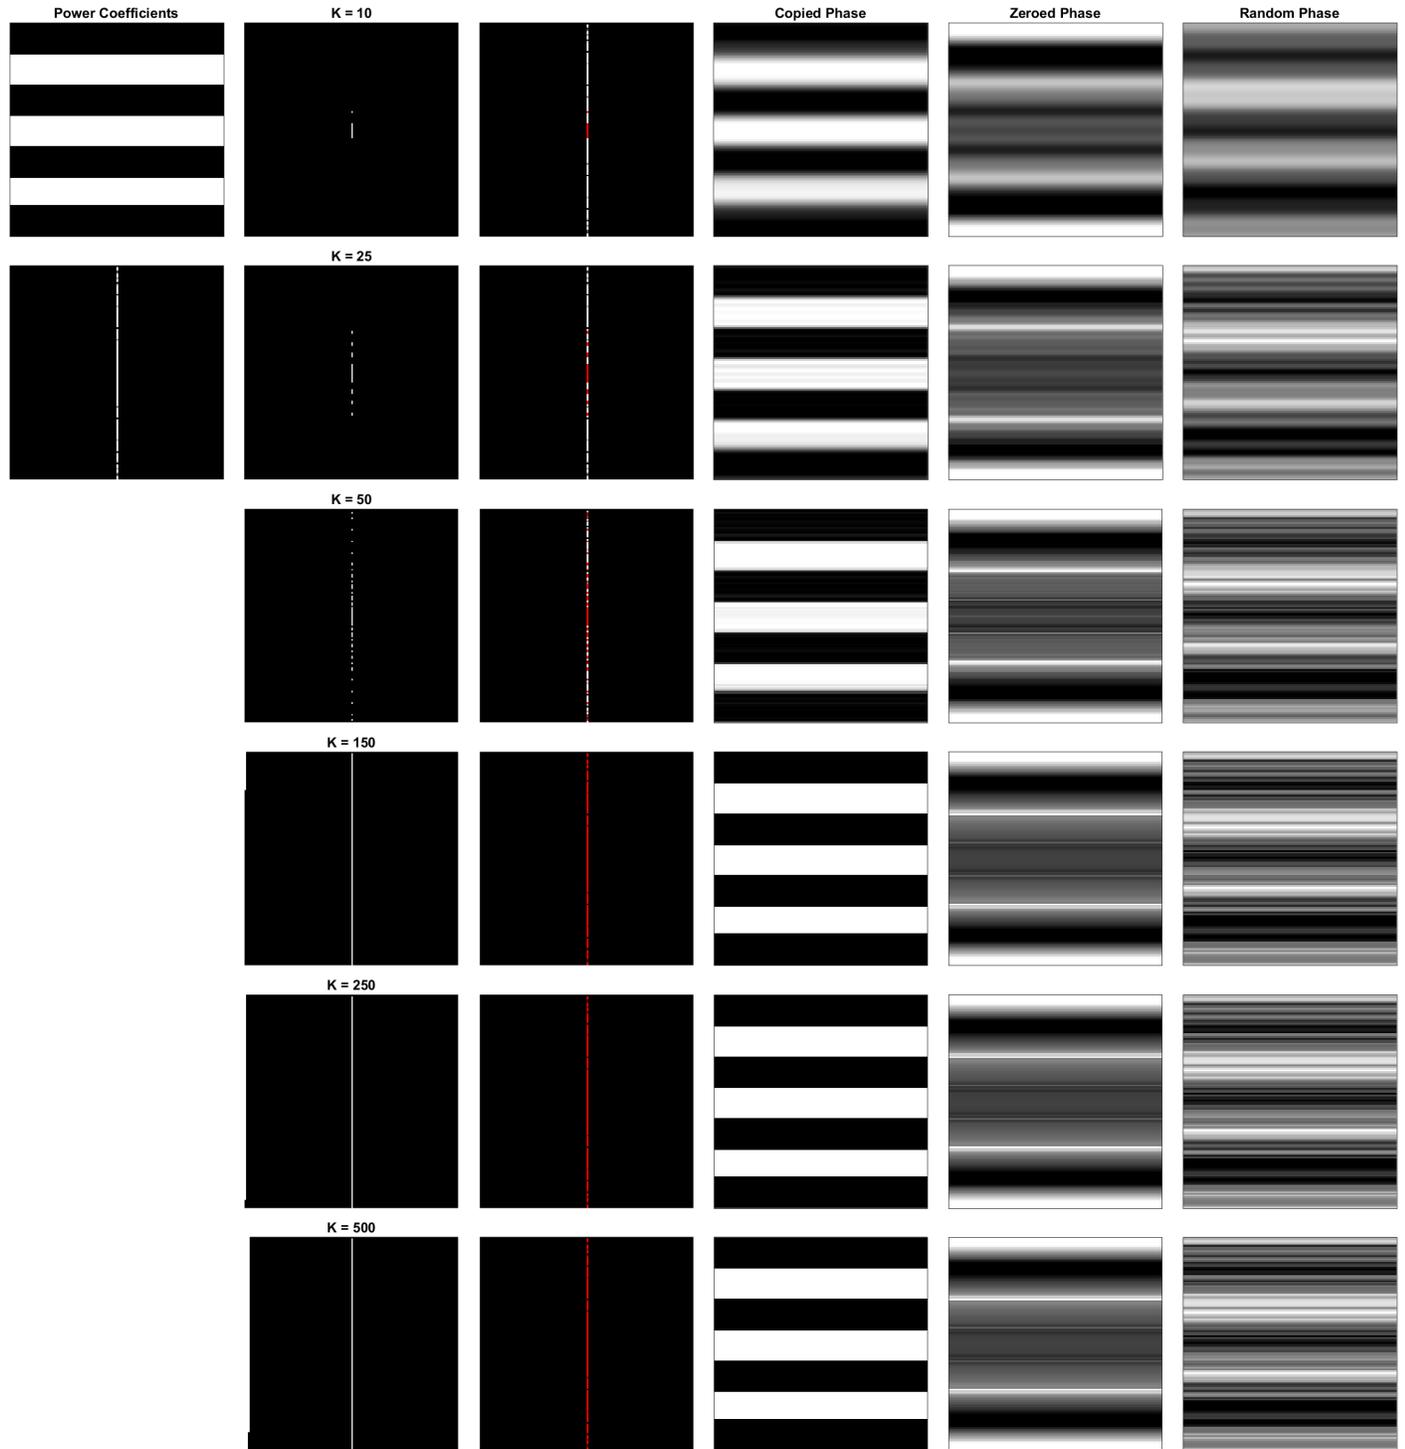
(h) Composition 04



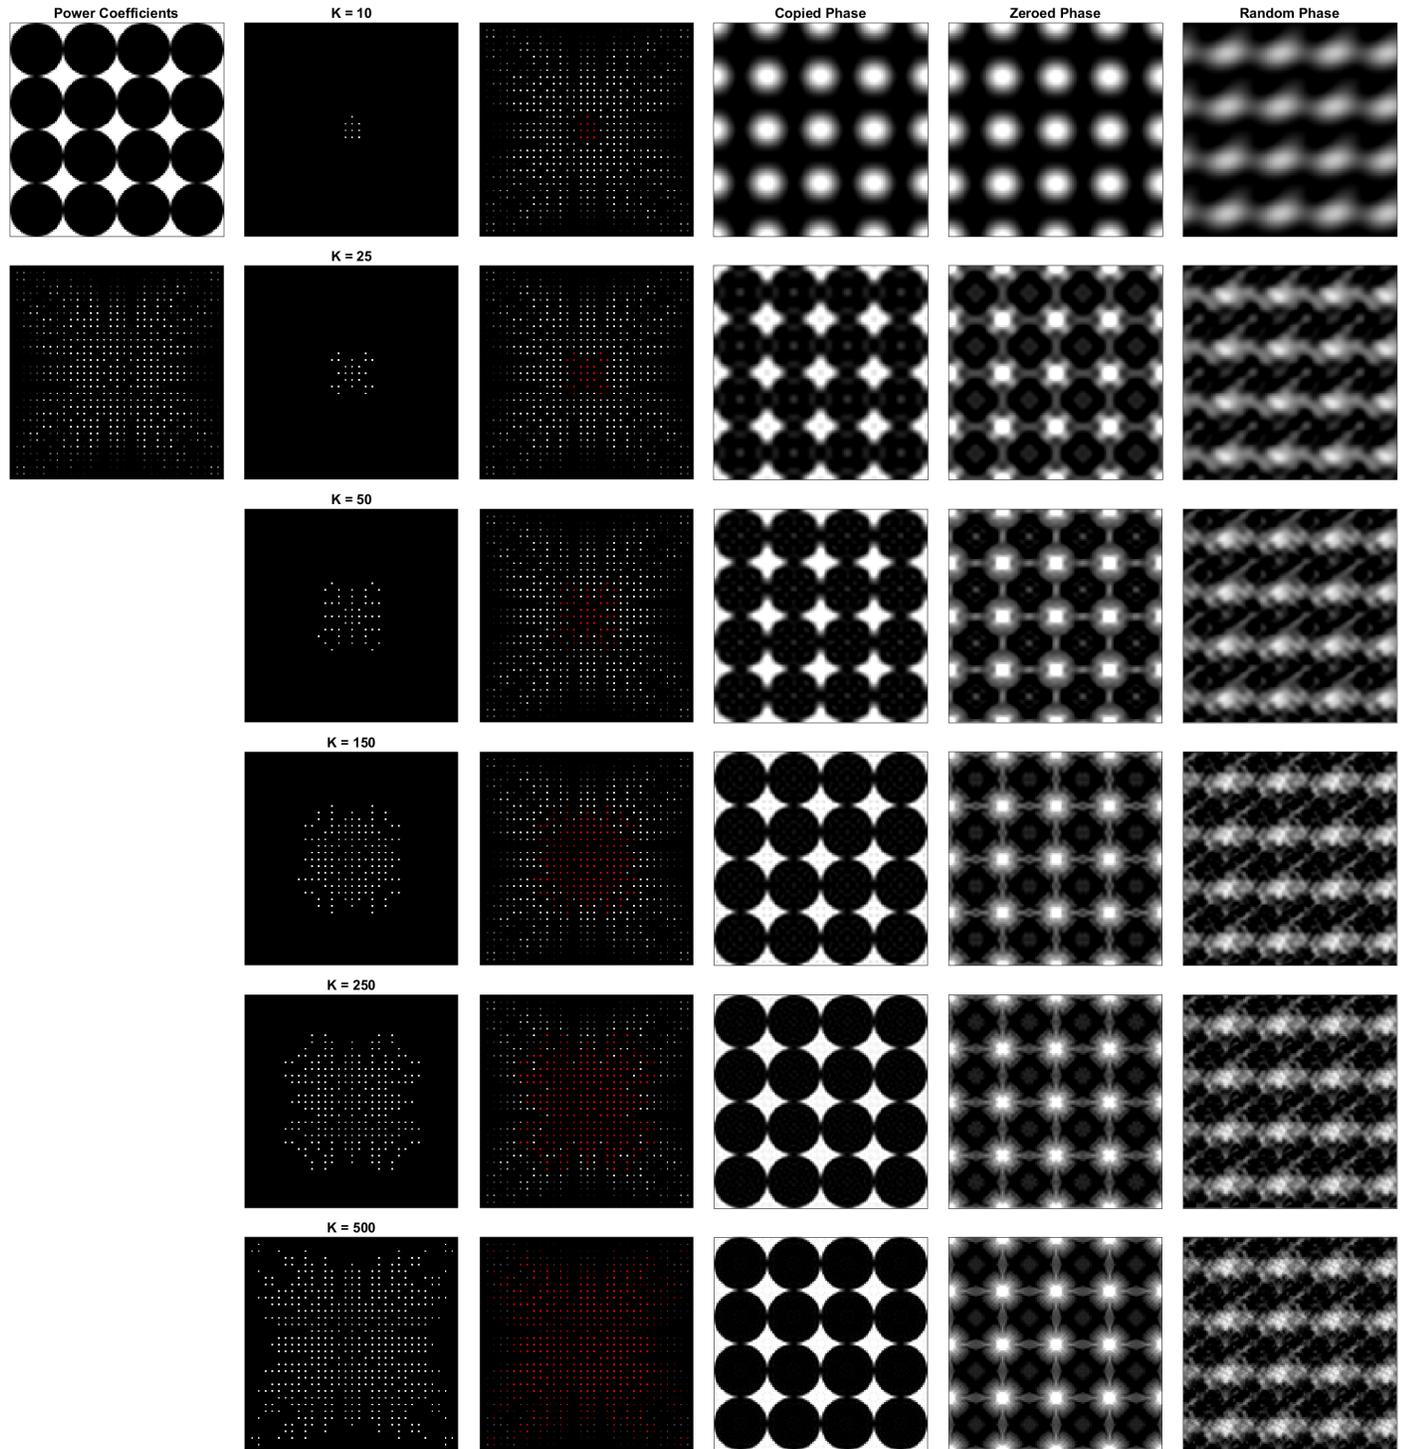
(i) Composition 05



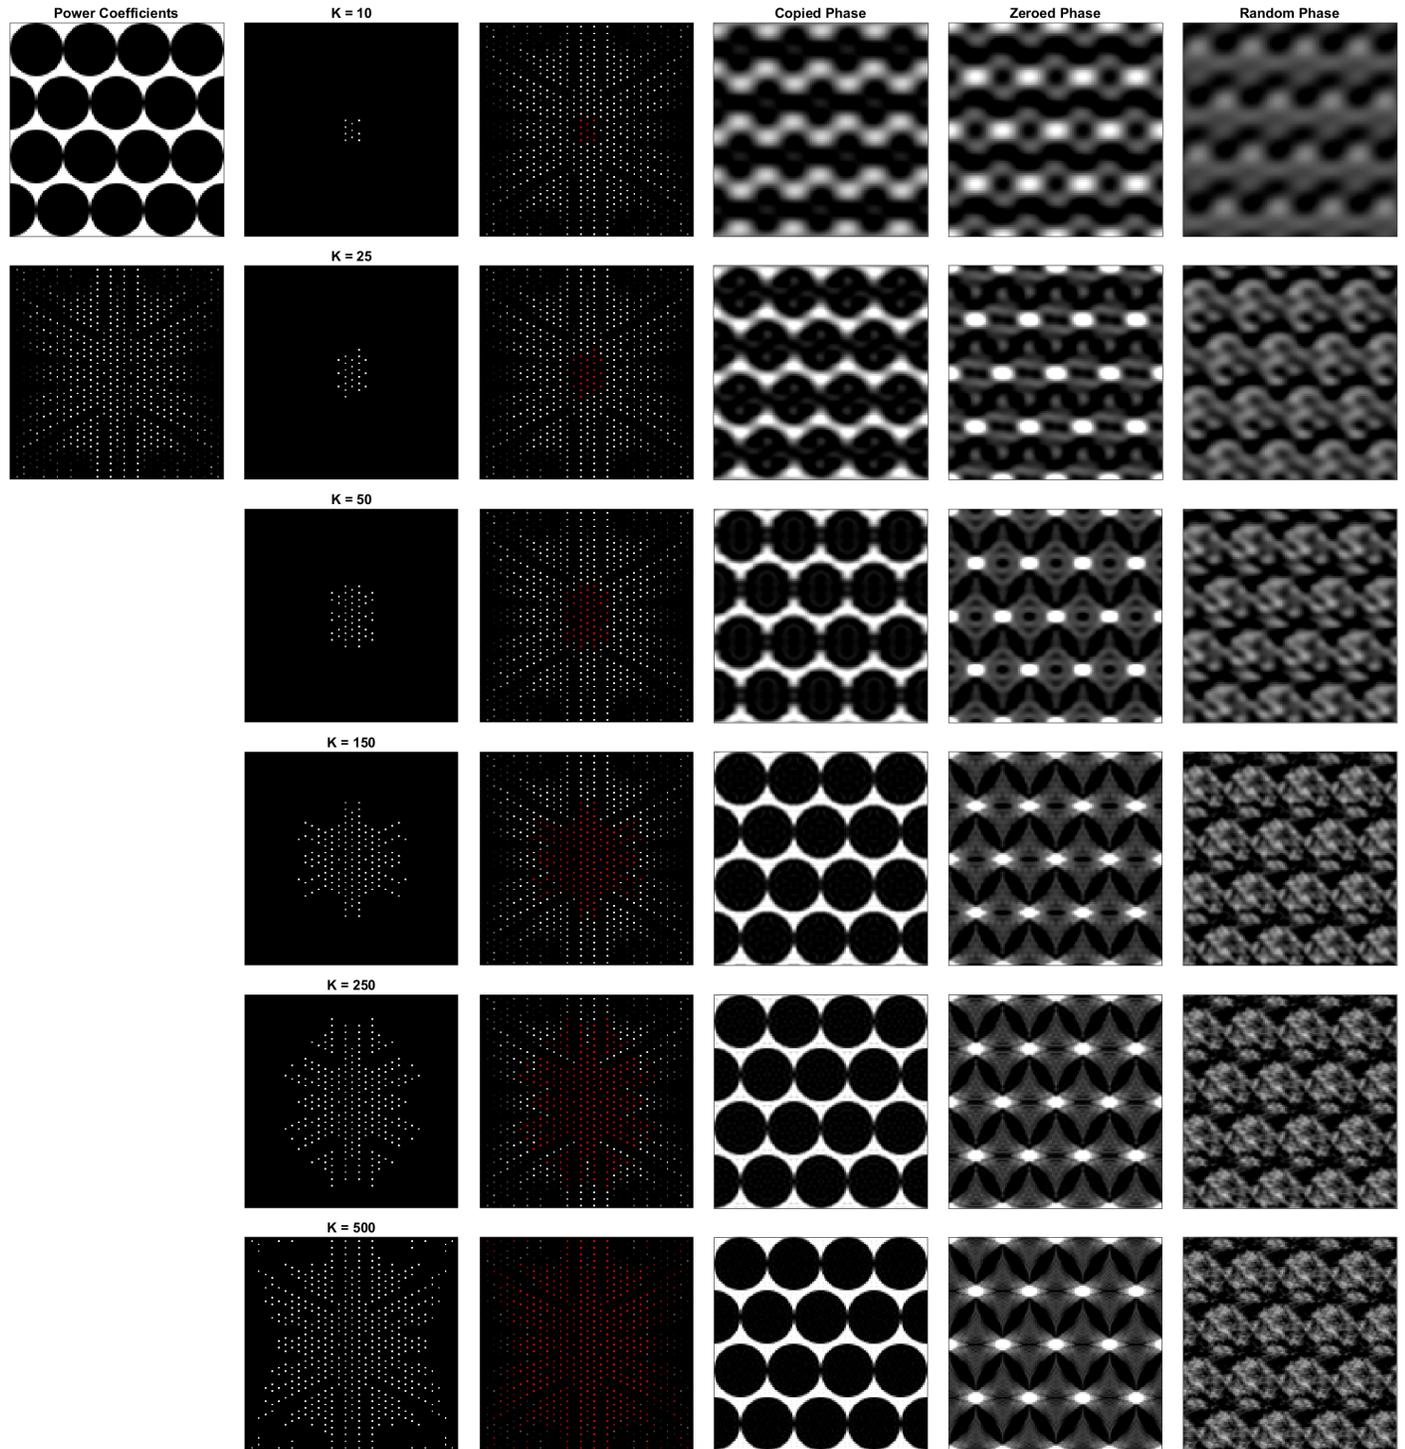
(j) Composition 06



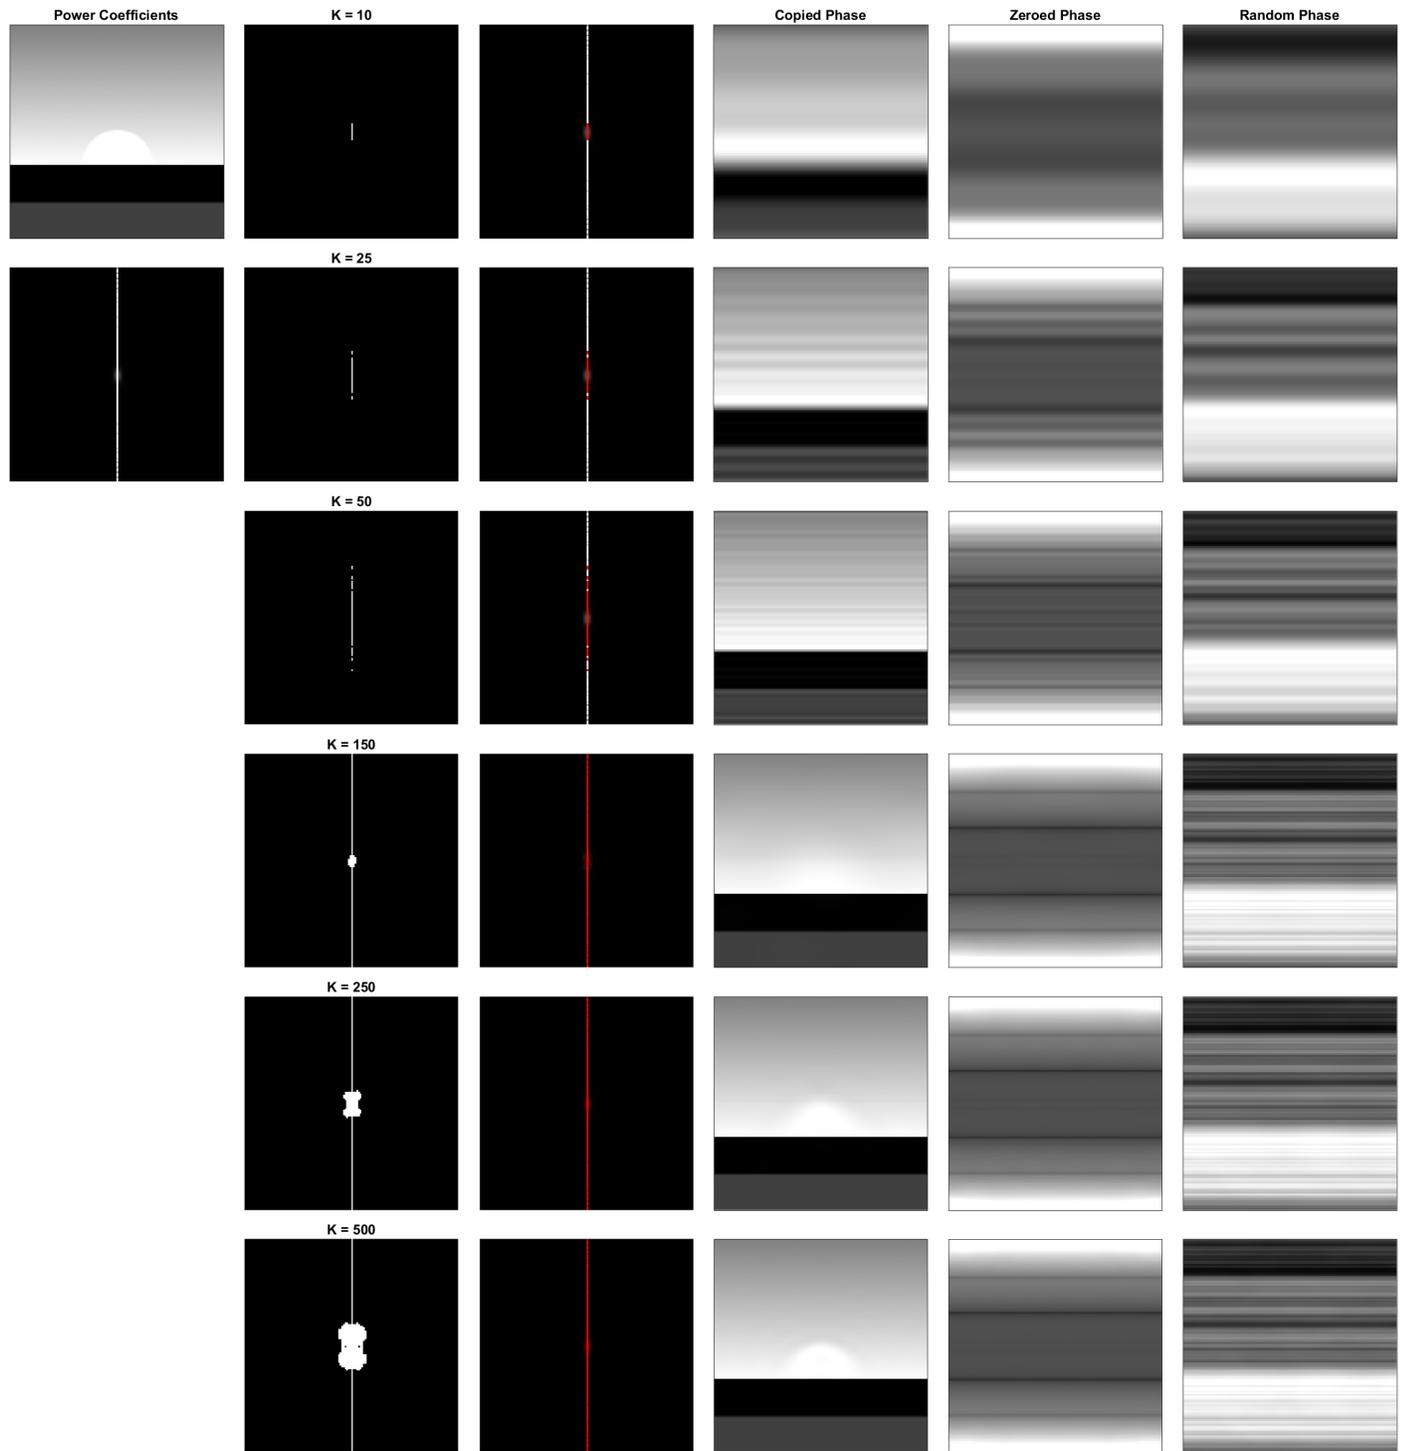
(k) Composition 07



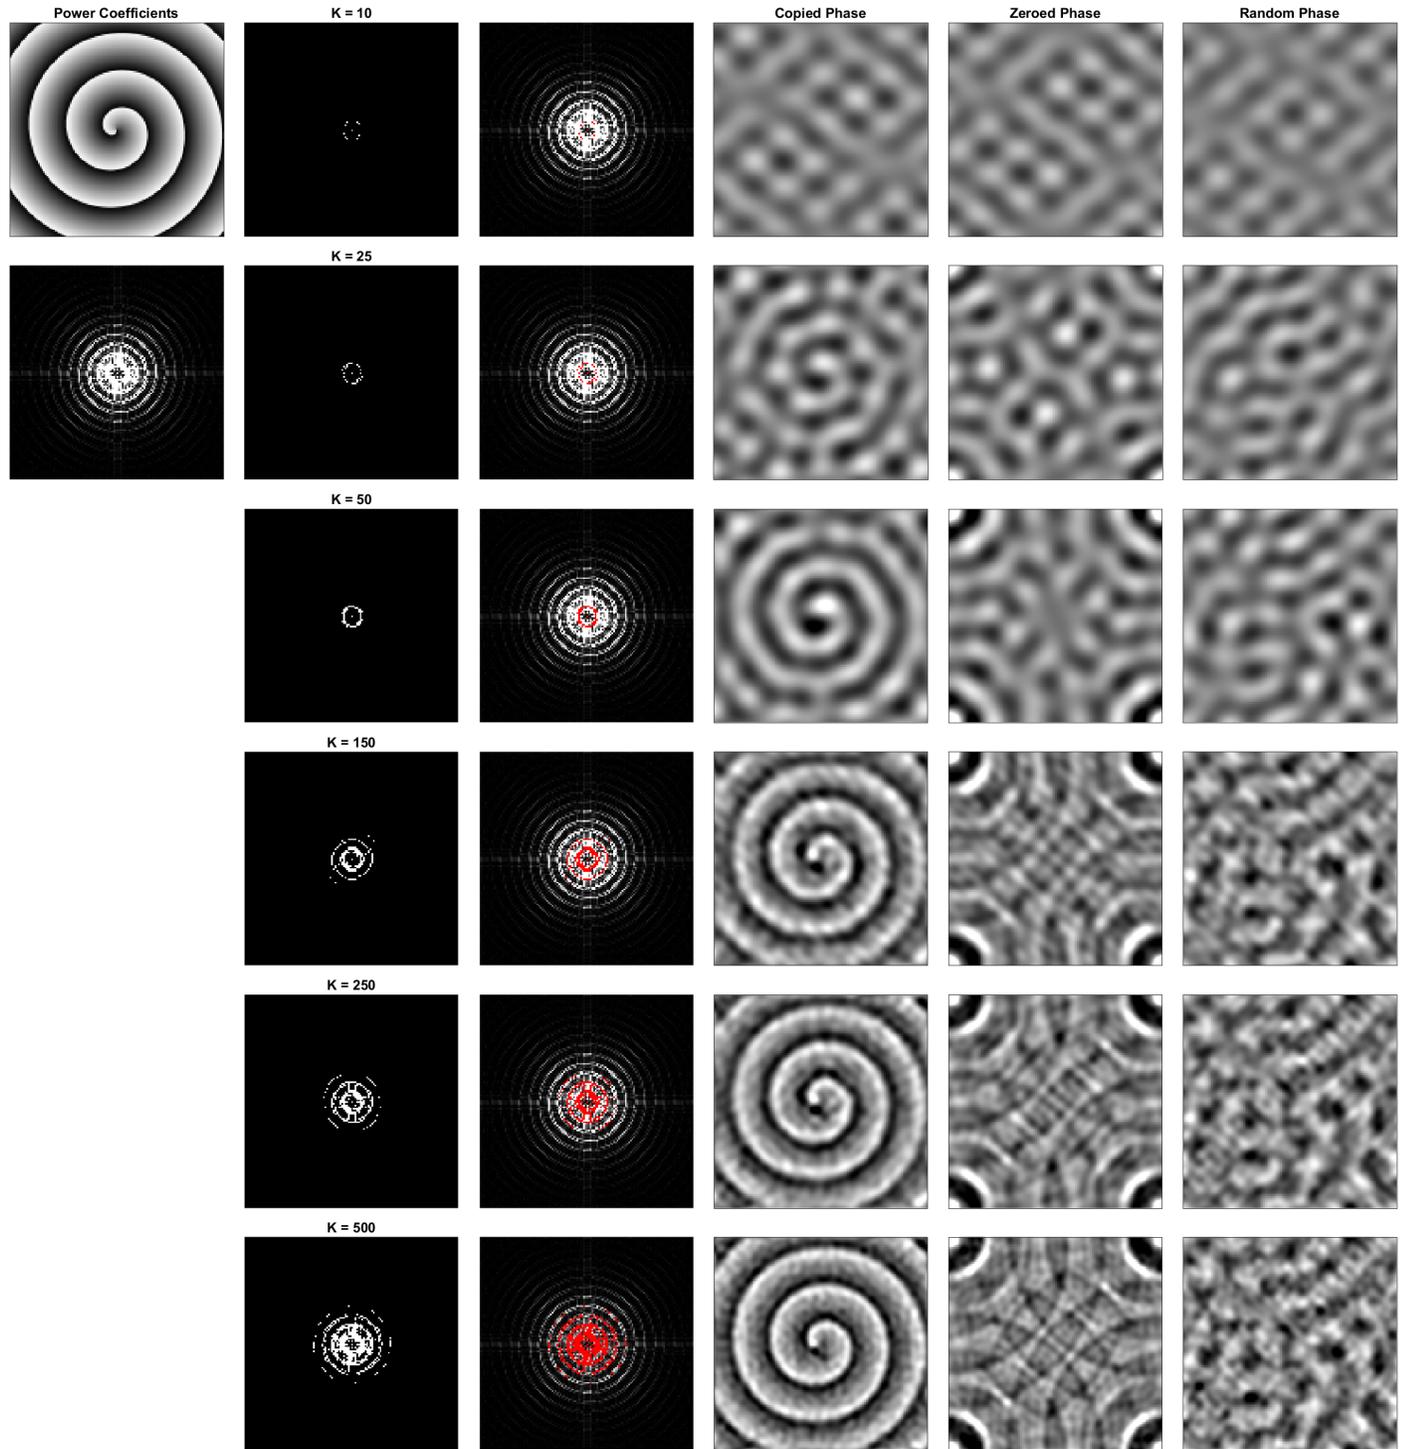
(1) Composition 08



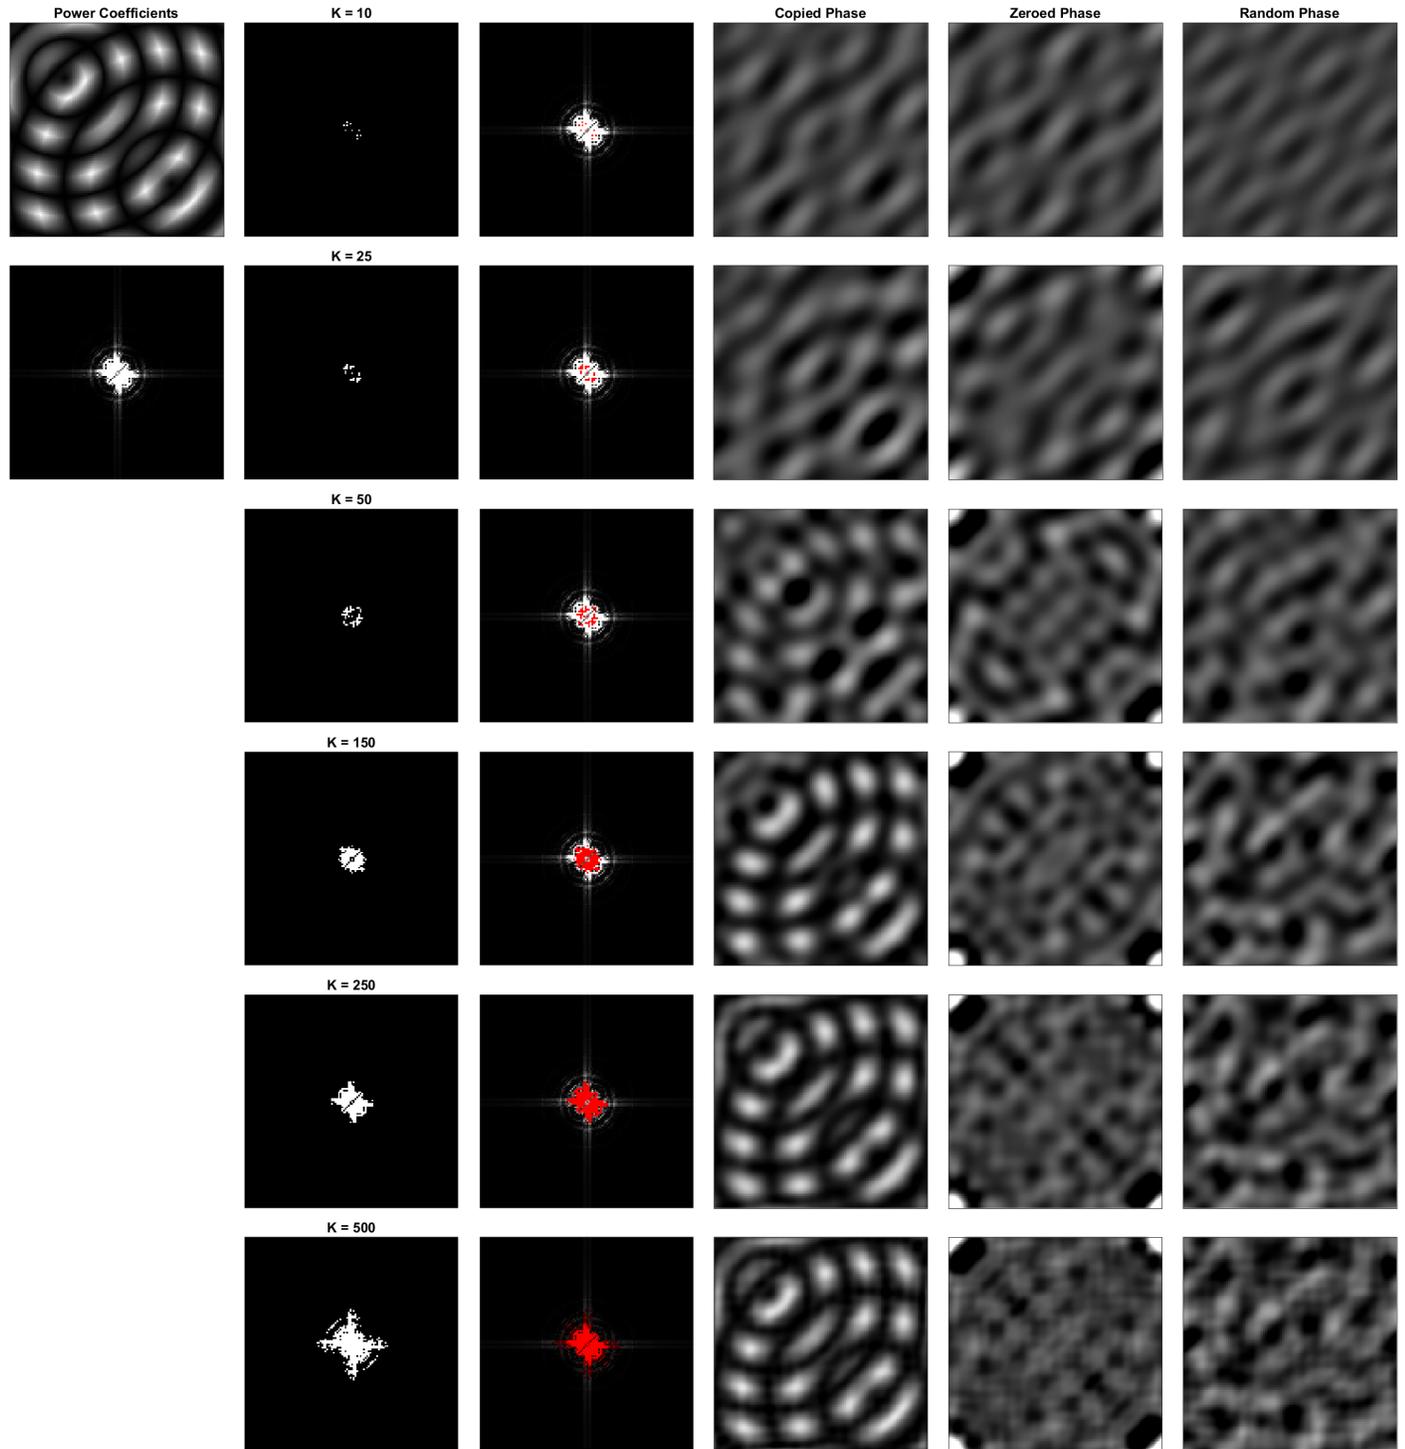
(m) Composition 09



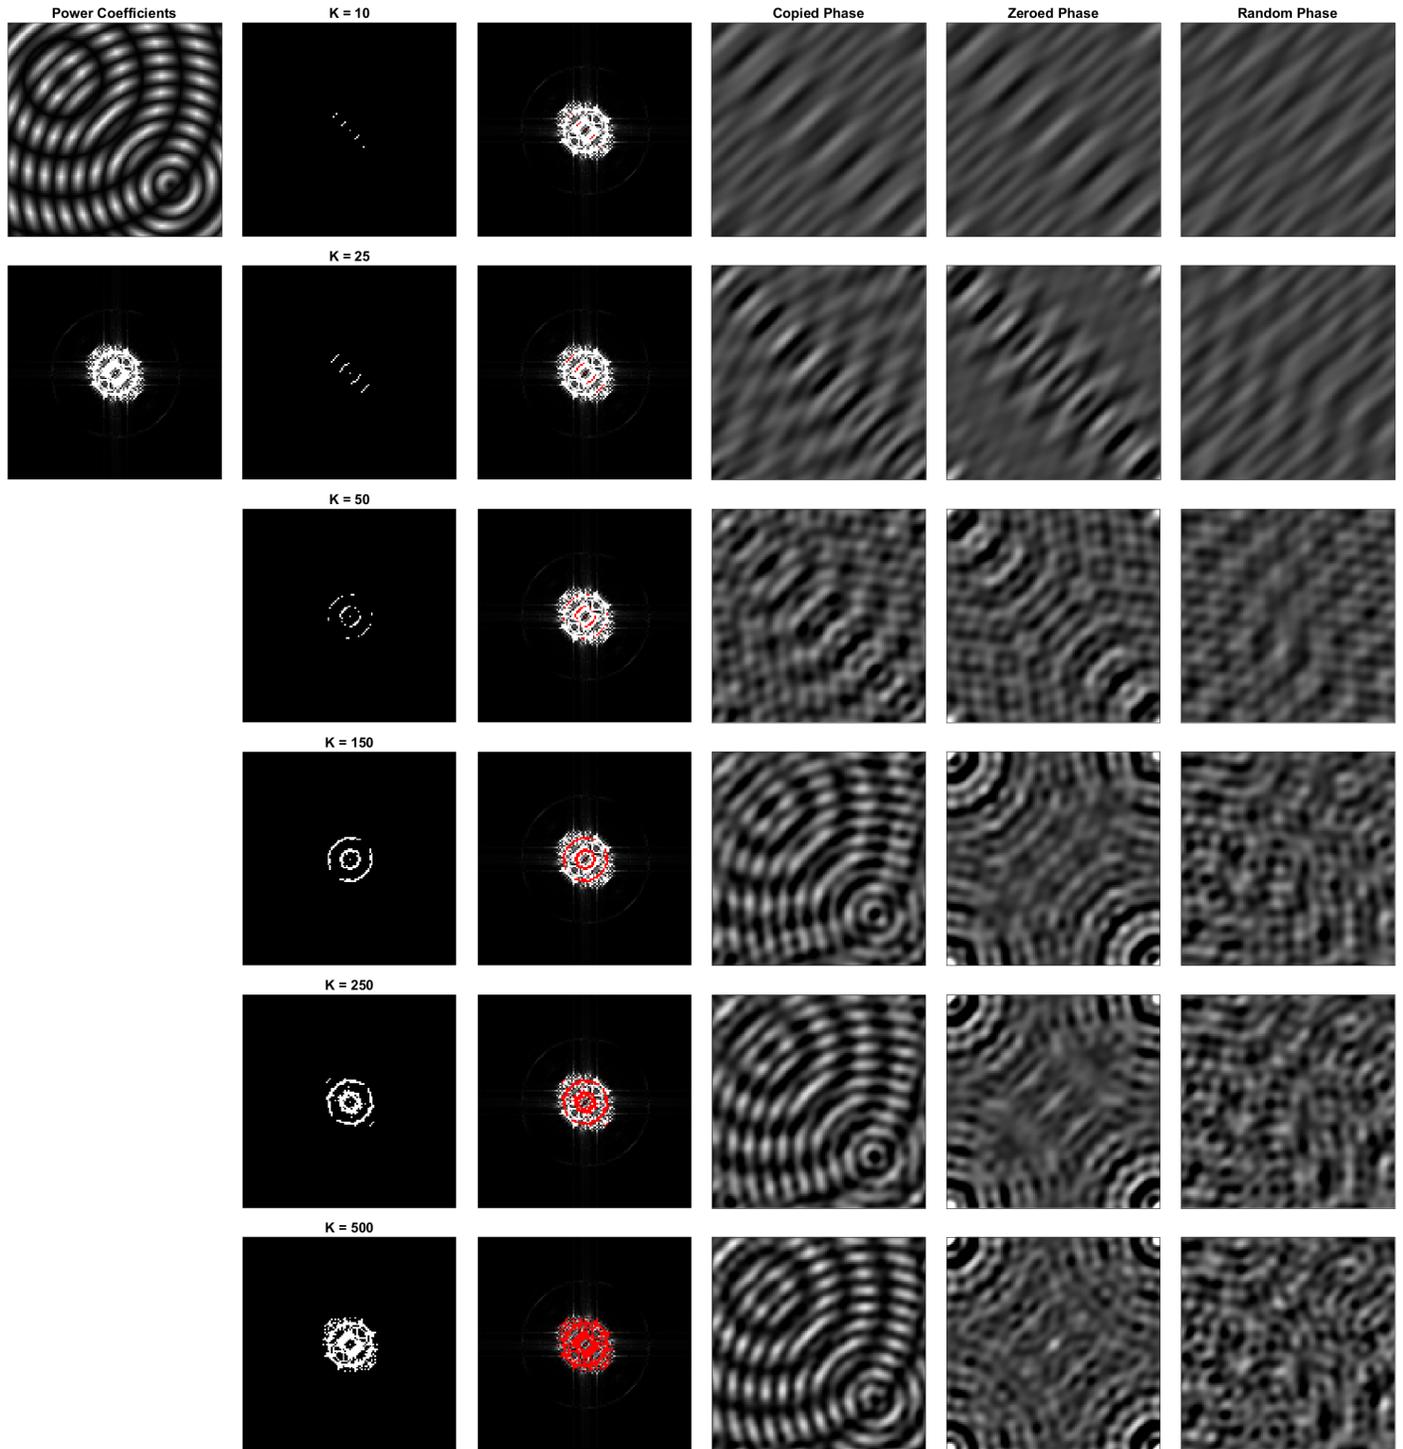
(n) Composition 10



(o) Composition 11



(p) Composition 12



# Appendix B

## Extended Results Data

Figure B.1: Experiment J3 Compositional Summary Charts & Examples;  $K = 10$

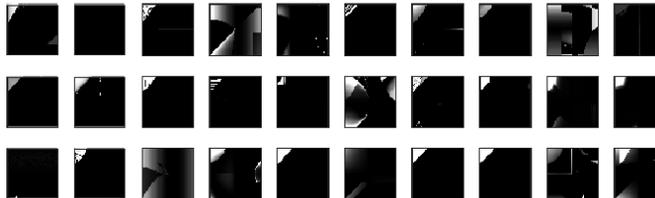
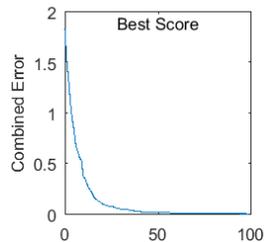
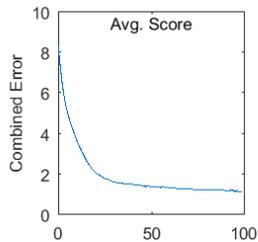
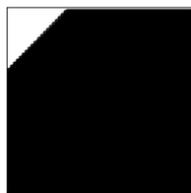
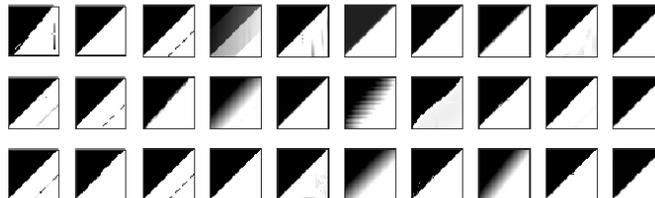
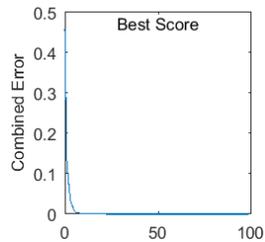
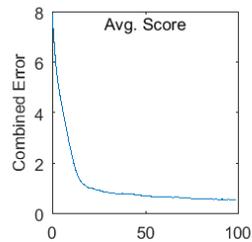
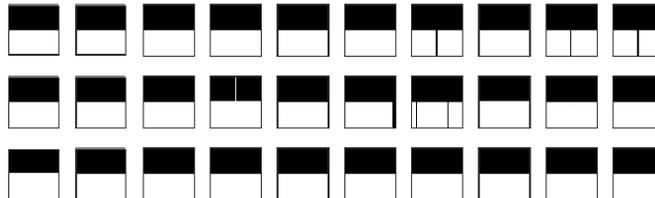
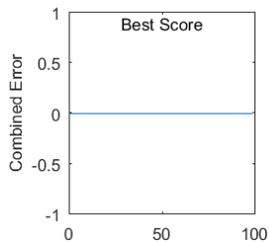
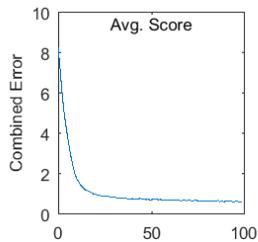
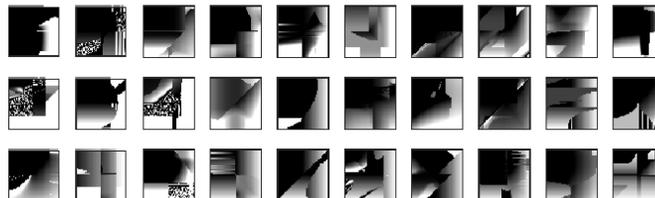
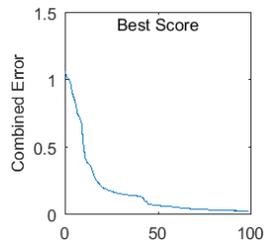
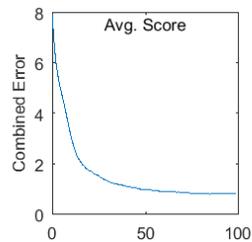
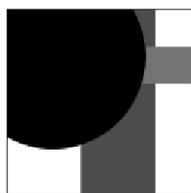


Figure B.1: Experiment J3 Compositional Summary Charts & Examples;  $K = 10$

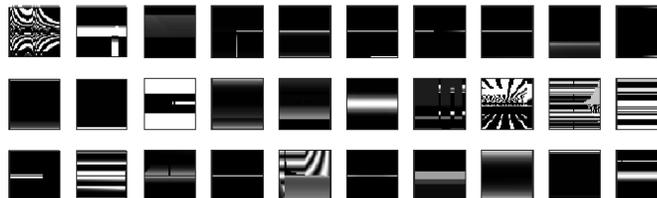
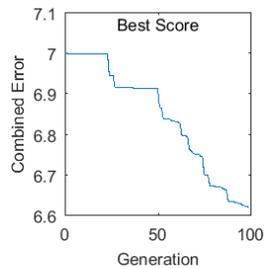
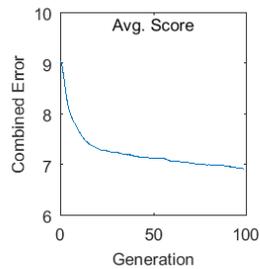
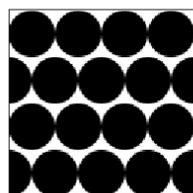
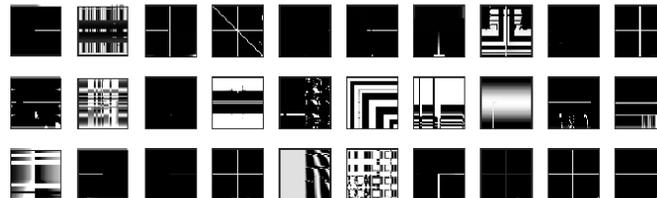
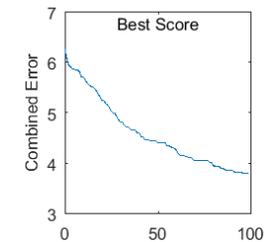
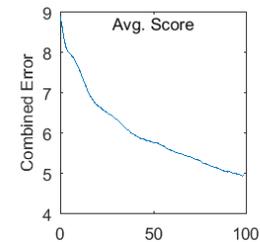
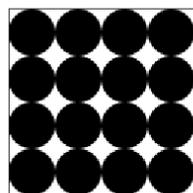
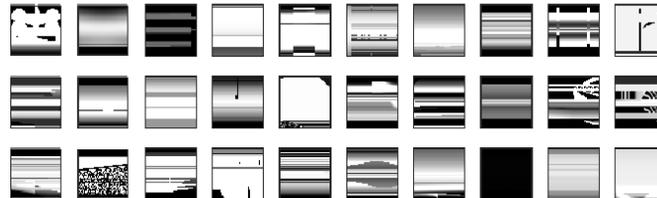
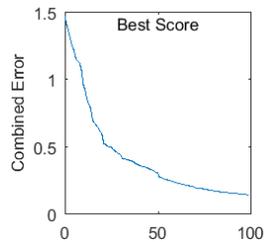
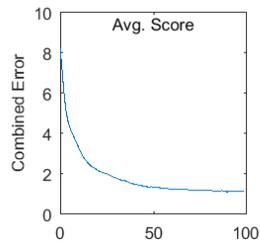
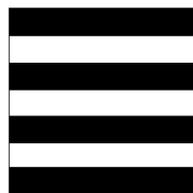
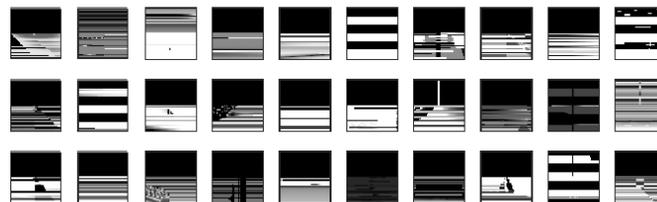
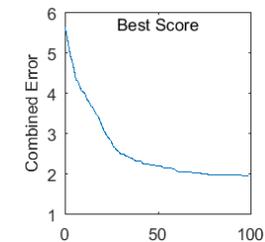
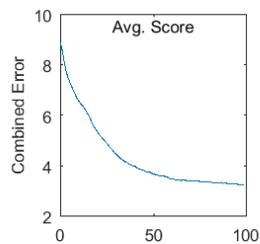
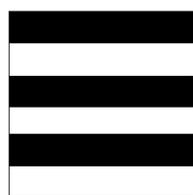


Figure B.1: Experiment J3 Compositional Summary Charts & Examples;  $K = 10$

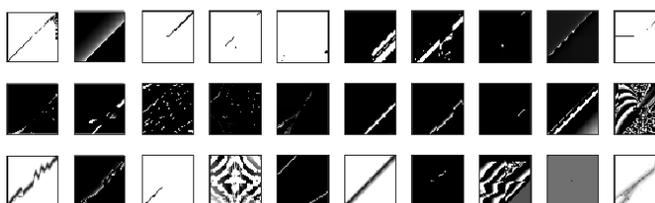
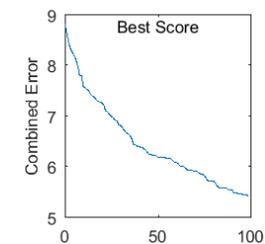
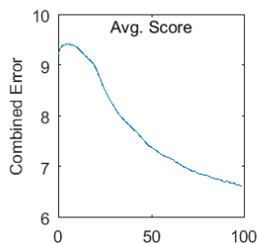
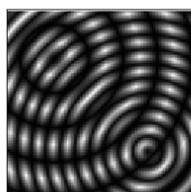
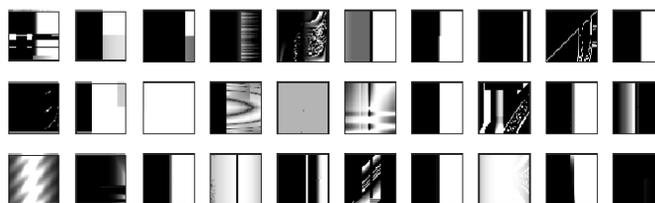
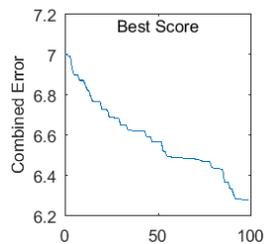
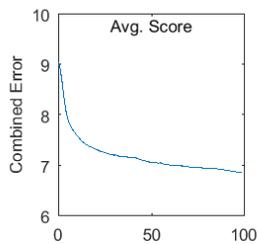
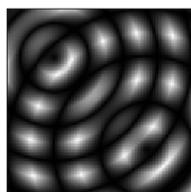
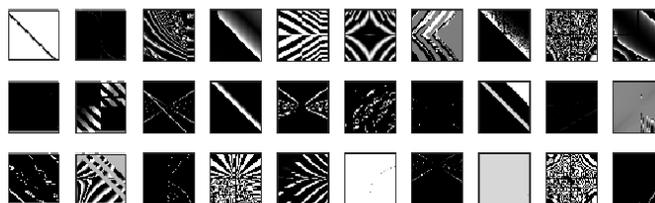
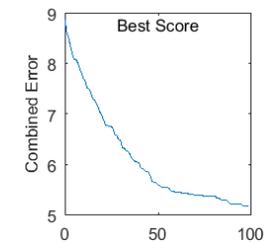
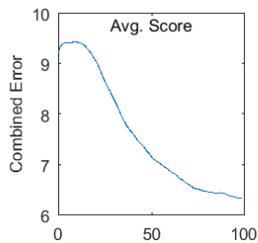
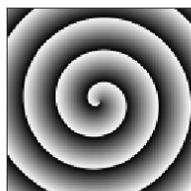
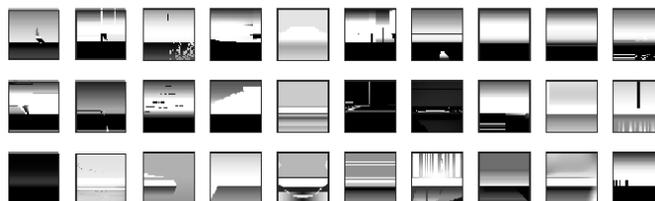
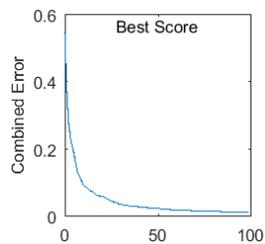
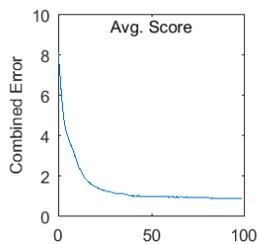


Figure B.2: Experiment K3 Compositional Summary Charts & Examples;  $K = 10$

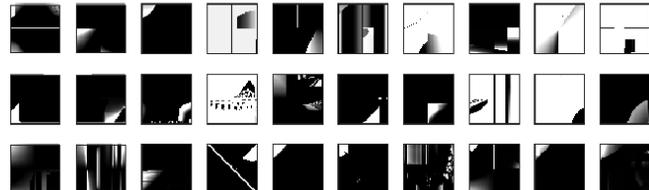
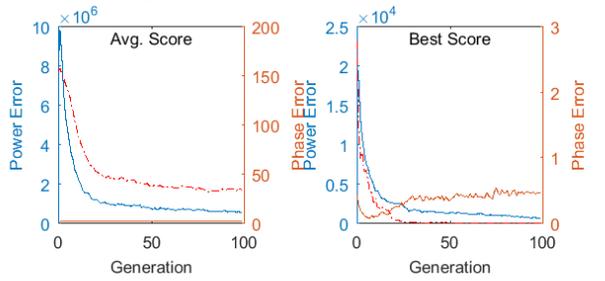
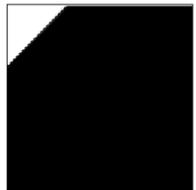
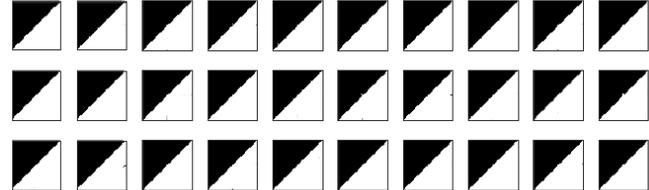
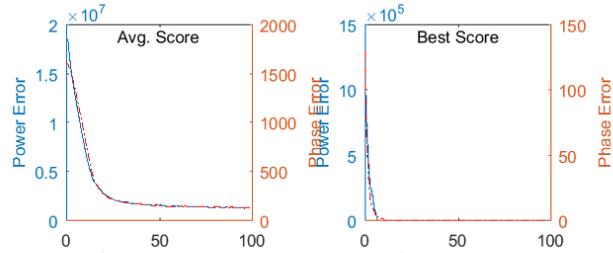
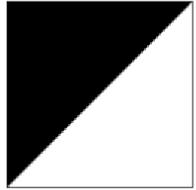
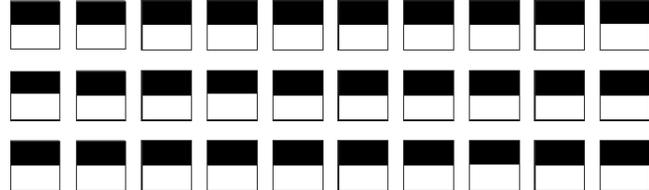
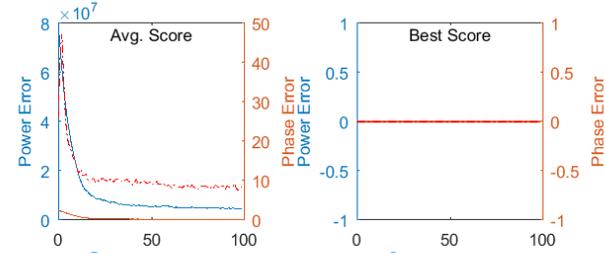
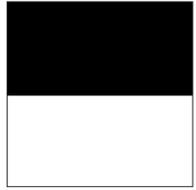
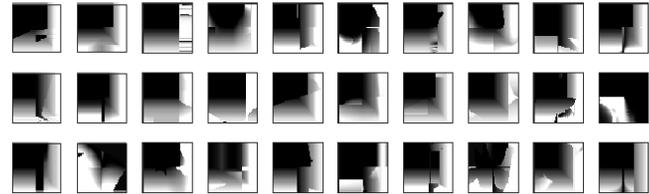
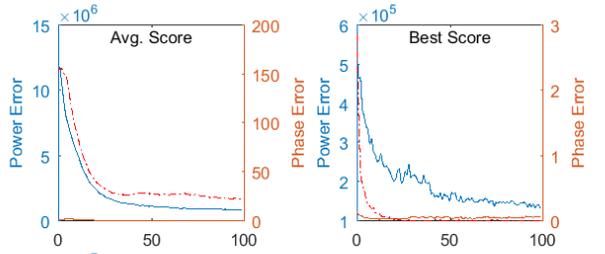


Figure B.2: Experiment K3 Compositional Summary Charts & Examples;  $K = 10$

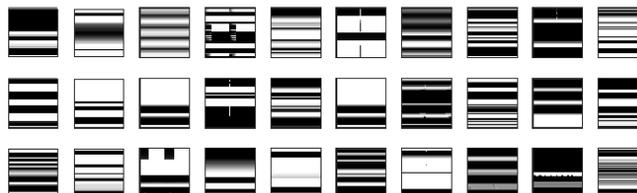
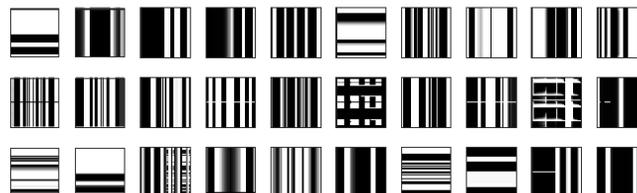
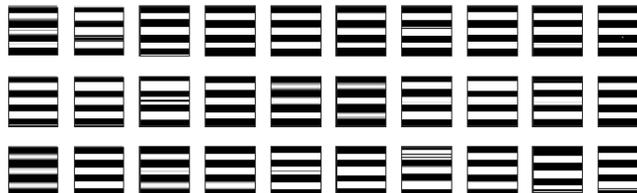
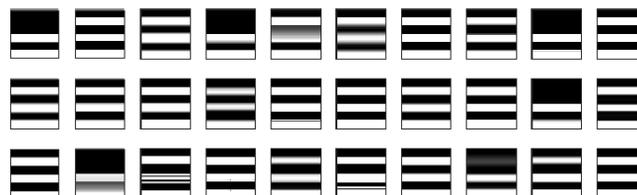
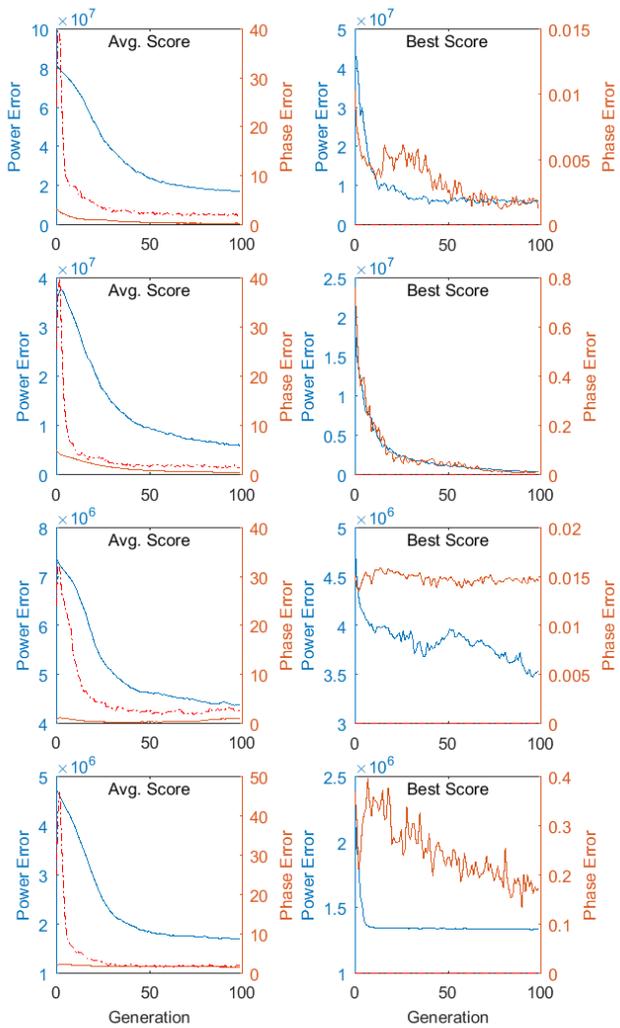
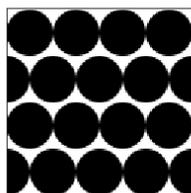
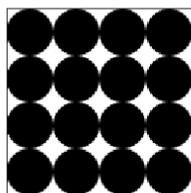


Figure B.2: Experiment K3 Compositional Summary Charts & Examples;  $K = 10$

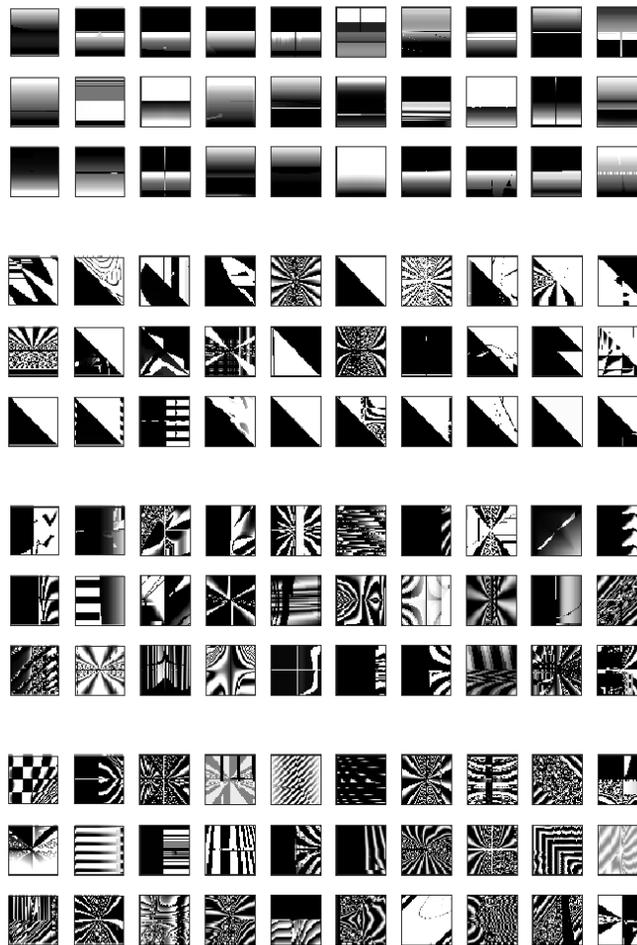
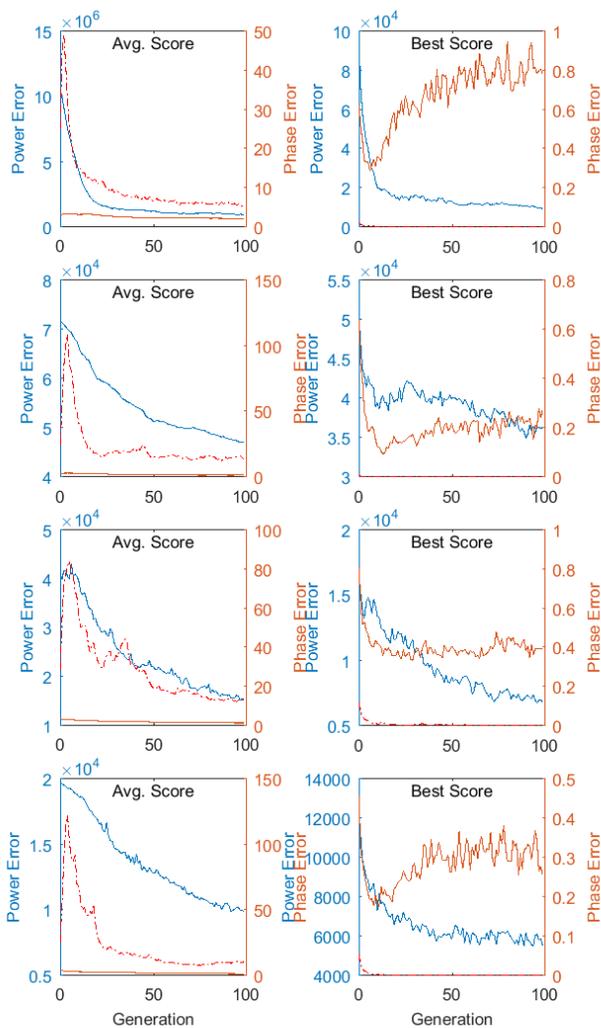
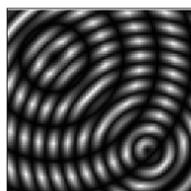
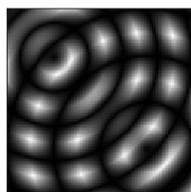
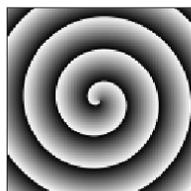


Figure B.3: Experiment P1 Compositional Summary Charts & Examples;  $K = 10$

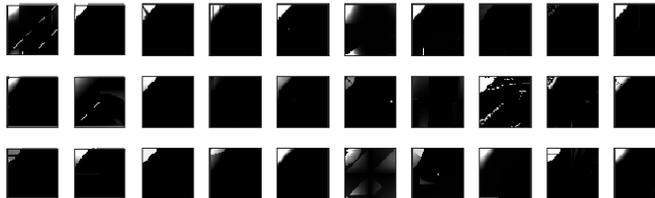
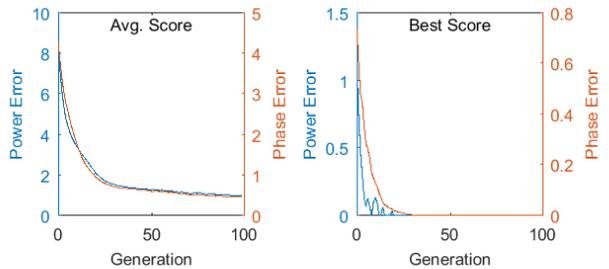
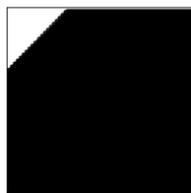
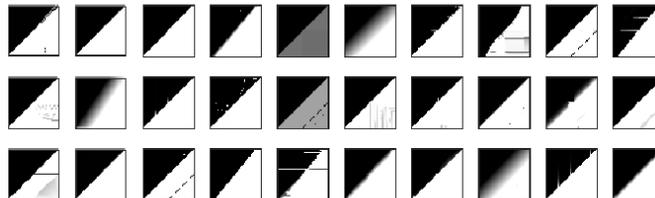
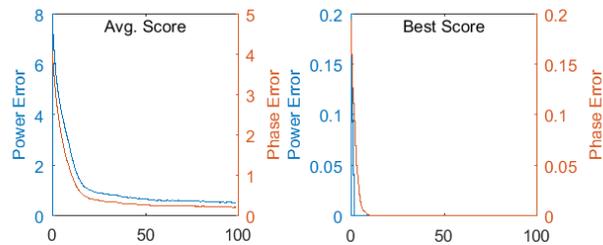
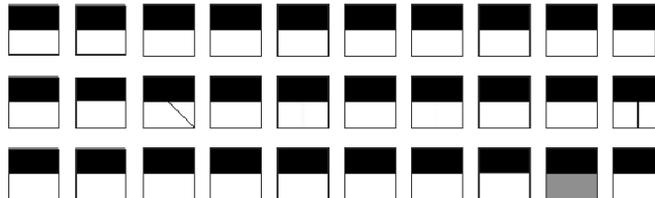
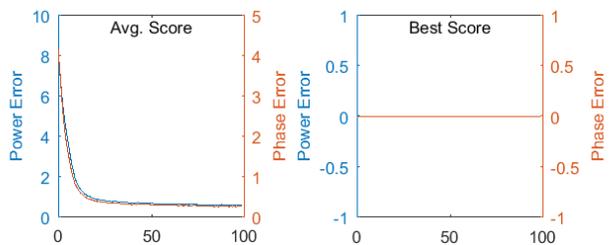
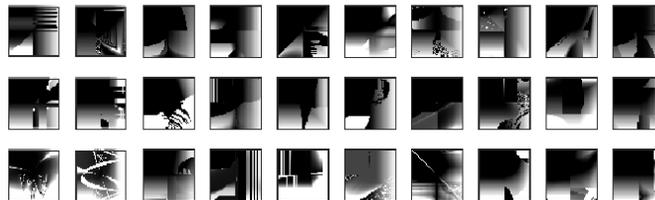
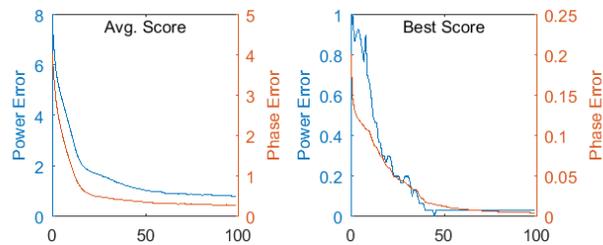
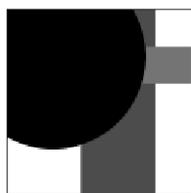


Figure B.3: Experiment P1 Compositional Summary Charts & Examples;  $K = 10$

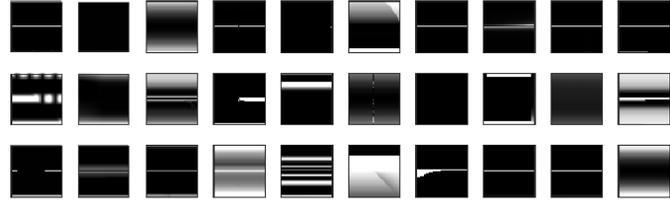
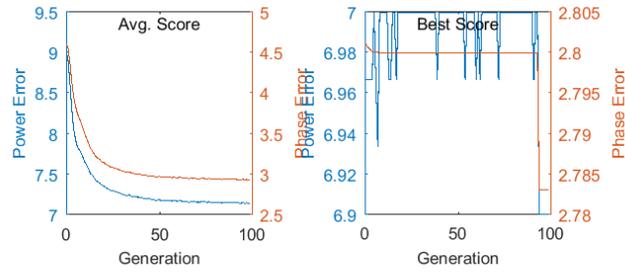
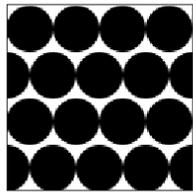
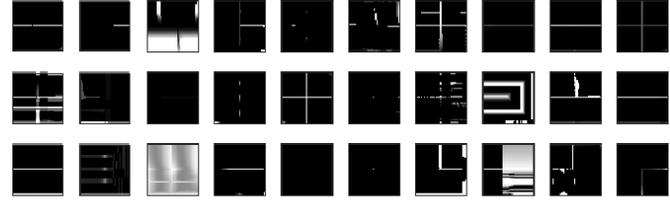
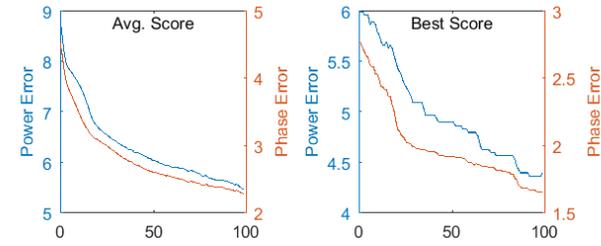
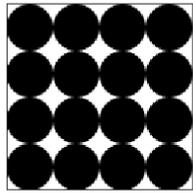
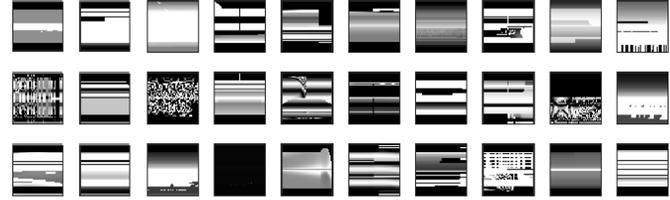
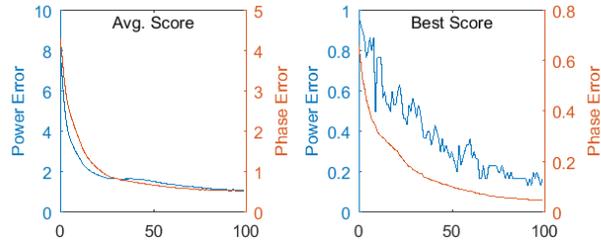
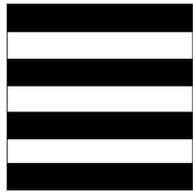
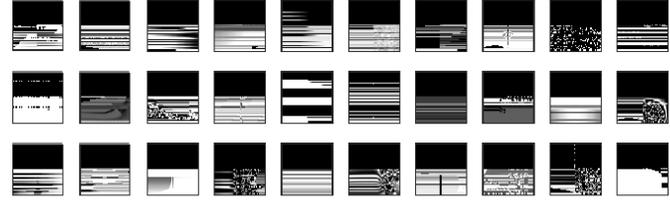
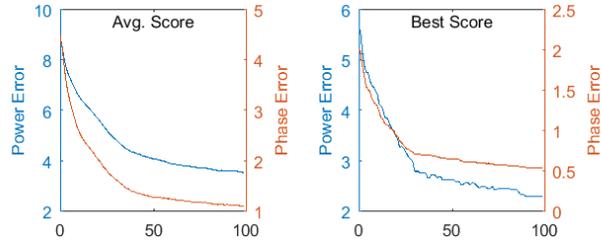
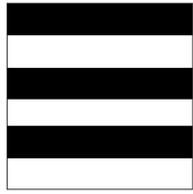


Figure B.3: Experiment P1 Compositional Summary Charts & Examples;  $K = 10$

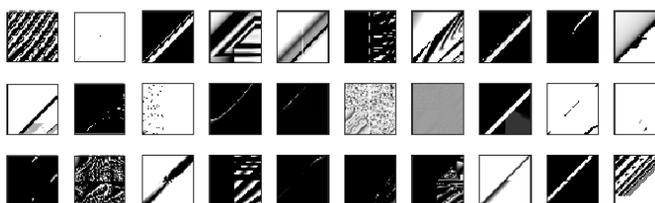
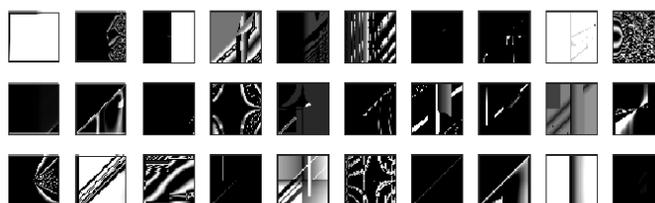
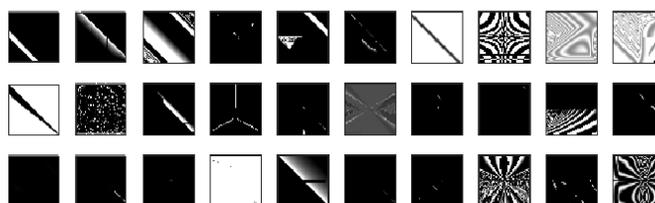
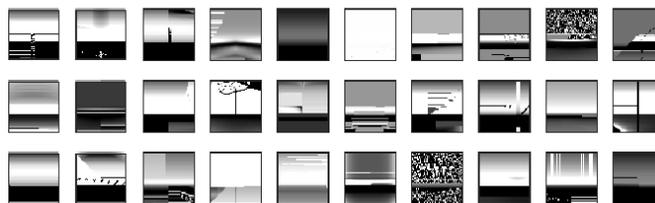
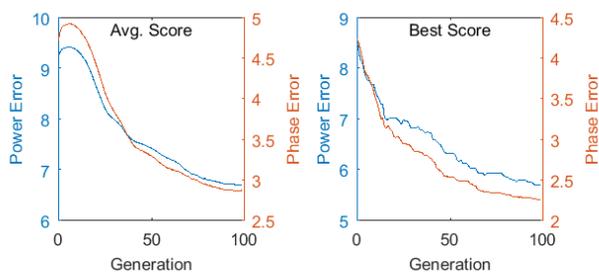
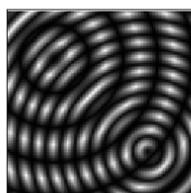
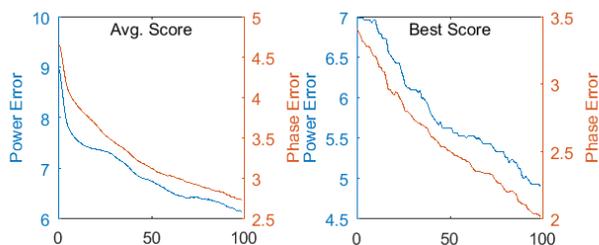
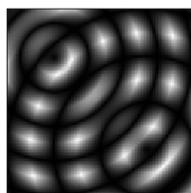
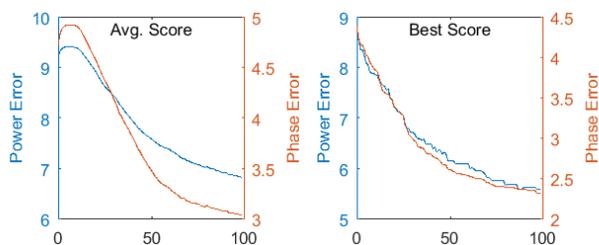
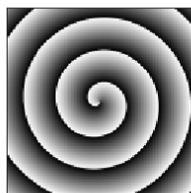
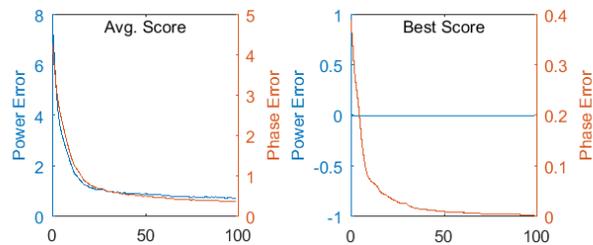


Figure B.4: Experiment P1 Compositional Summary Charts & Examples;  $K = 25$

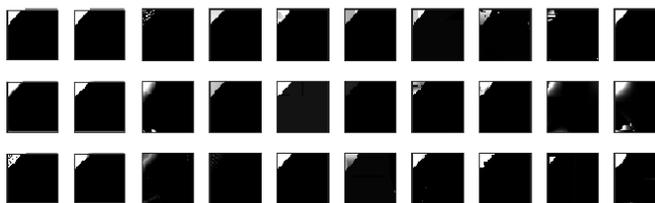
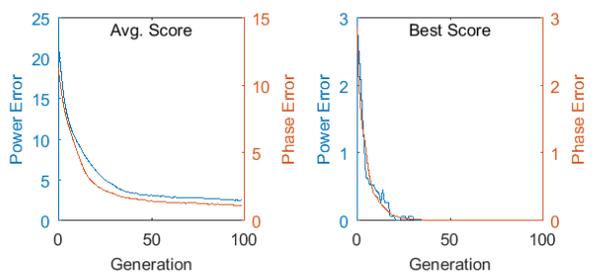
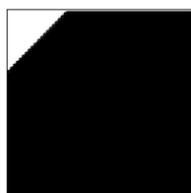
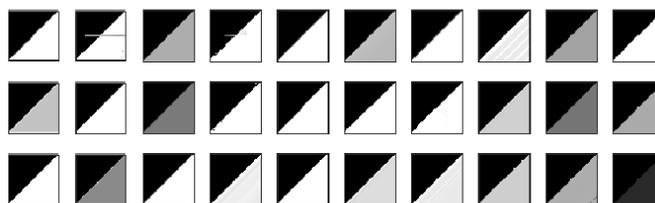
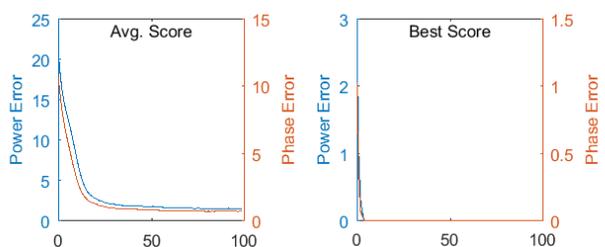
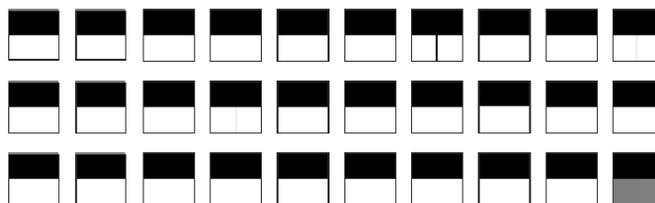
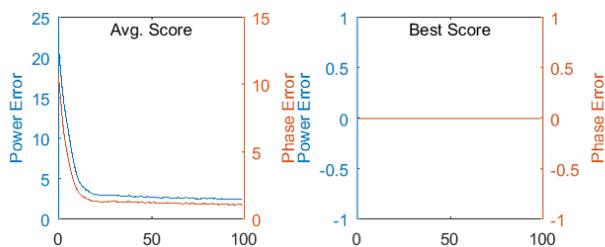
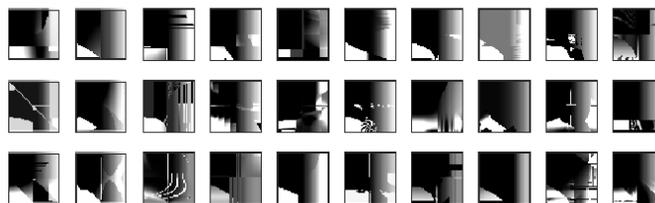
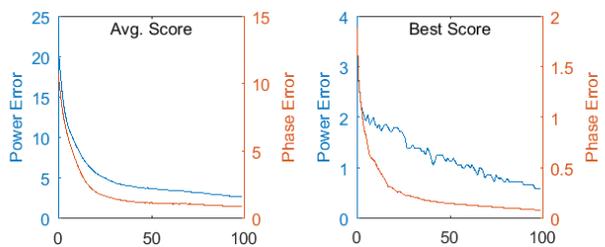
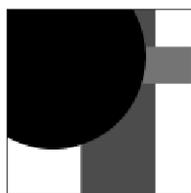


Figure B.4: Experiment P1 Compositional Summary Charts & Examples;  $K = 25$

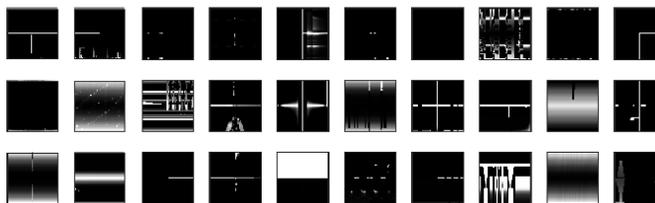
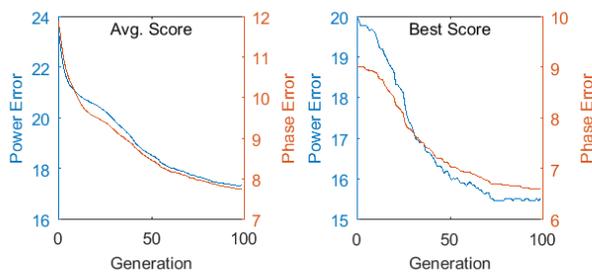
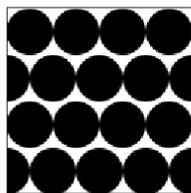
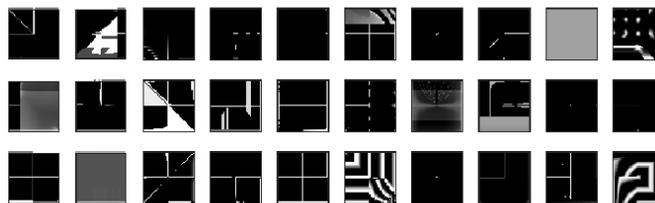
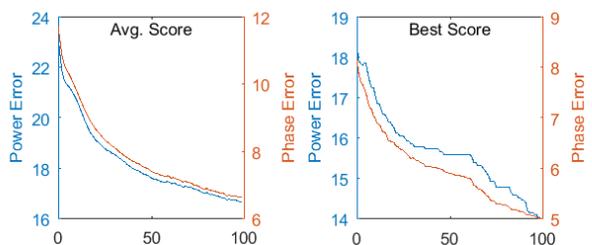
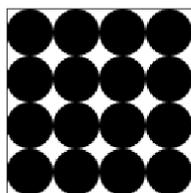
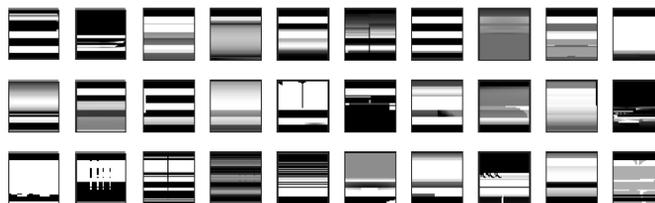
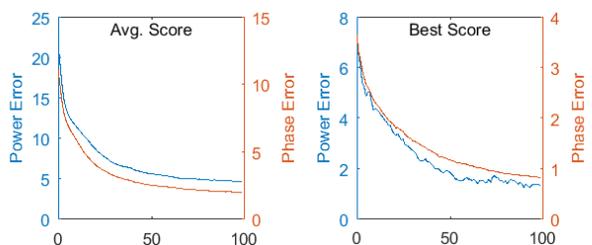
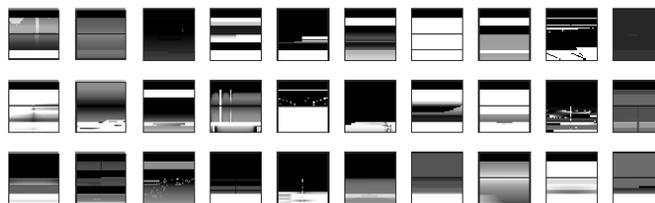
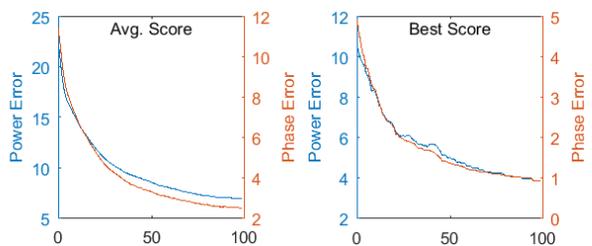


Figure B.4: Experiment P1 Compositional Summary Charts & Examples;  $K = 25$

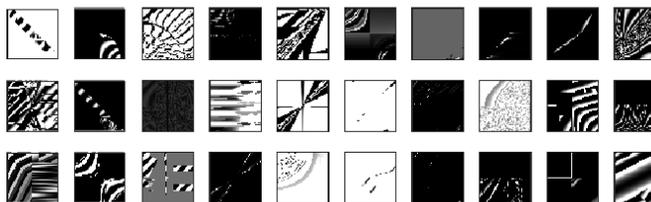
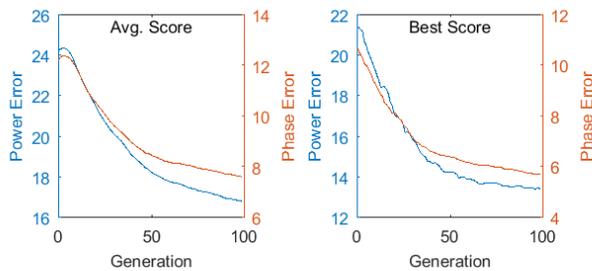
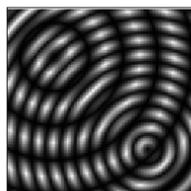
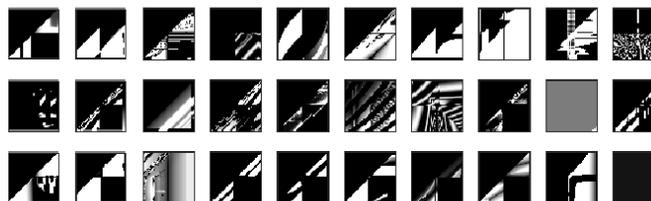
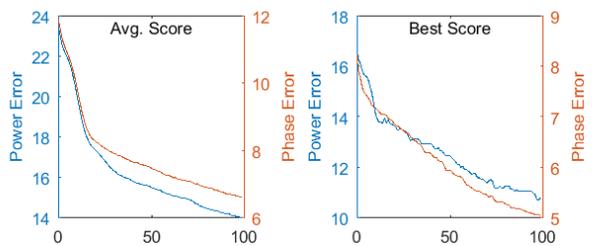
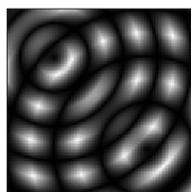
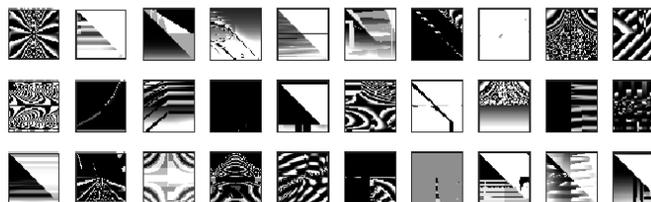
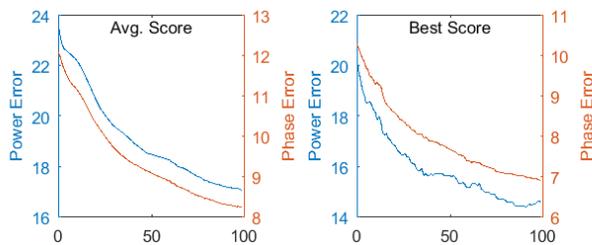
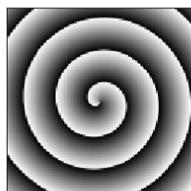
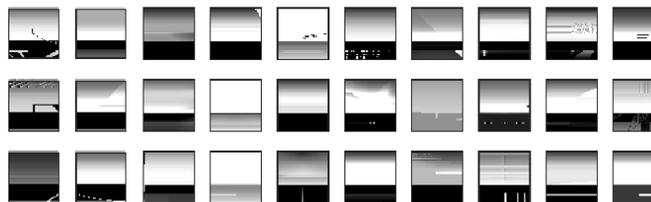
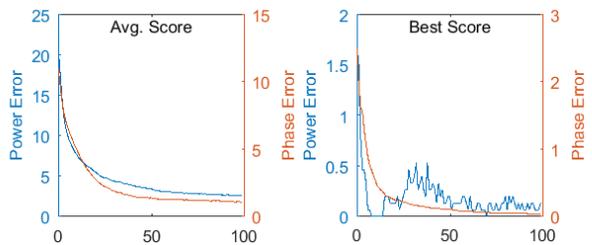


Figure B.5: Polar Coordinate Language Summary Charts & Examples;  $K = 10$

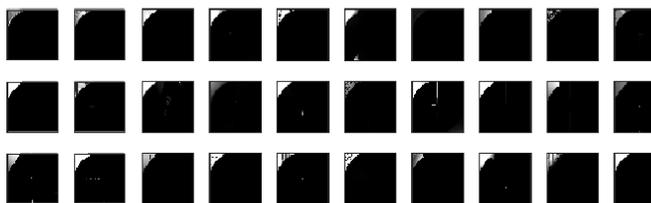
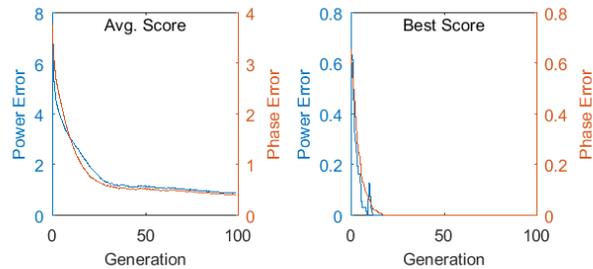
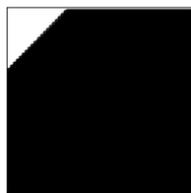
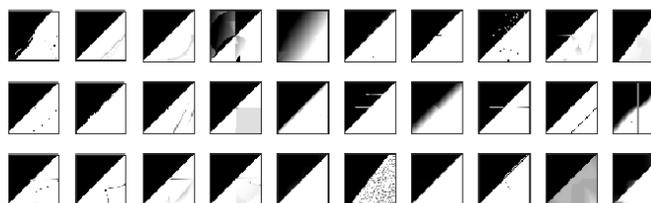
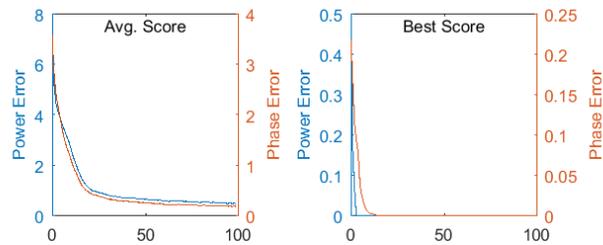
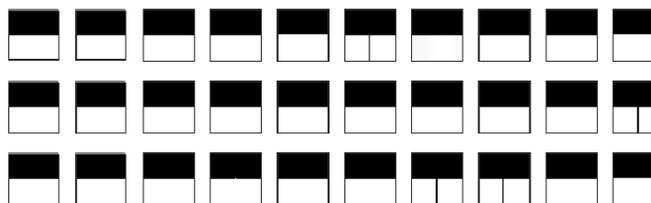
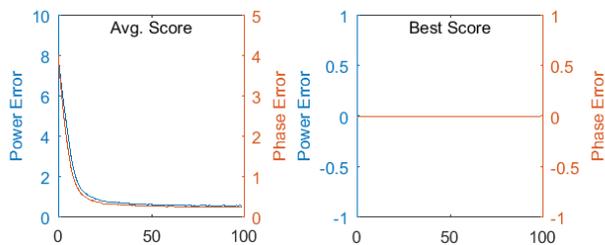
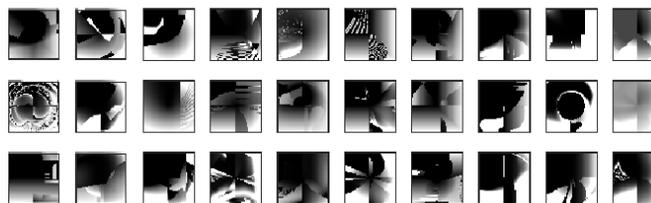
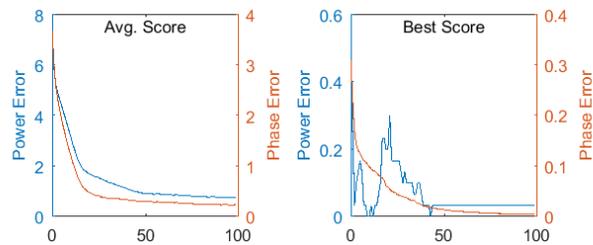
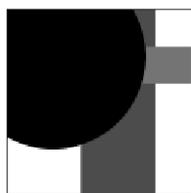


Figure B.5: Polar Coordinate Language Summary Charts & Examples;  $K = 10$

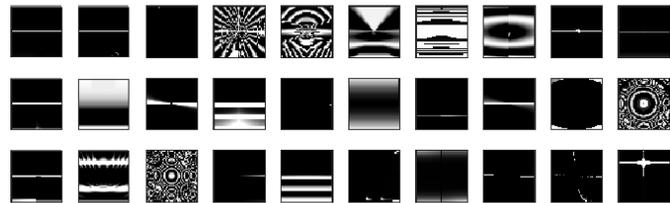
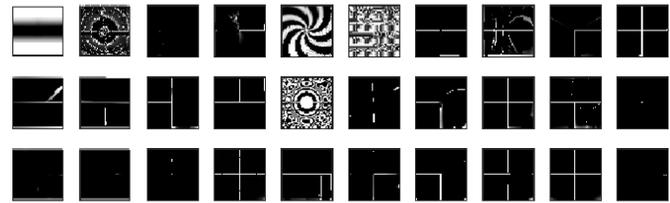
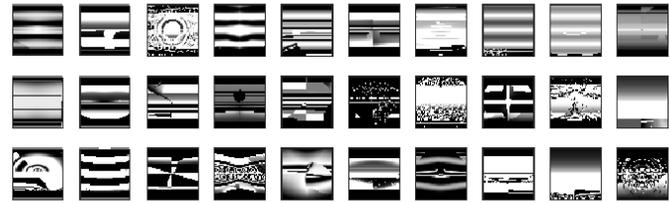
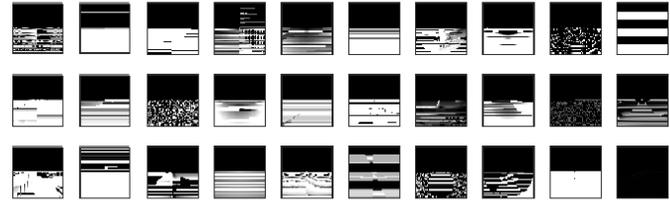
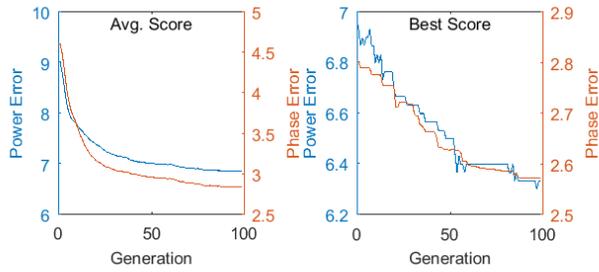
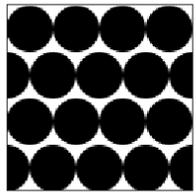
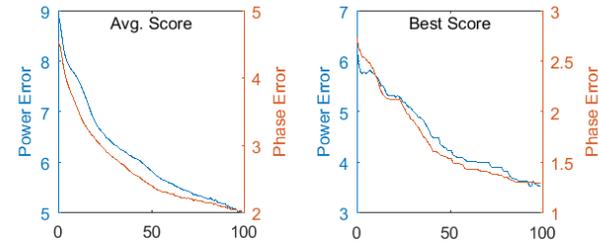
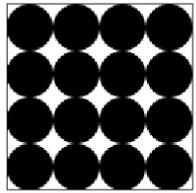
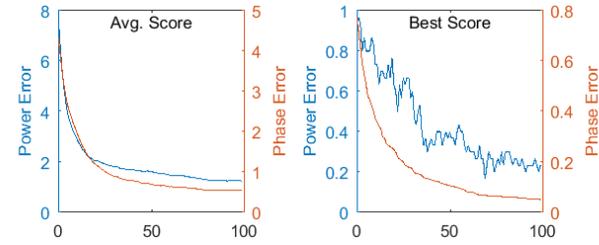
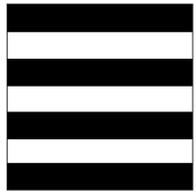
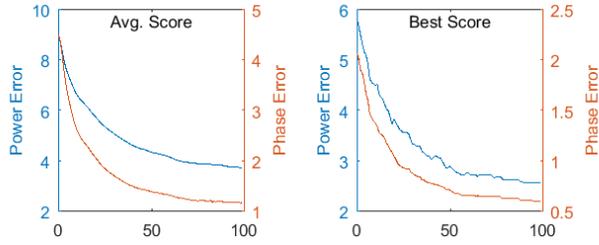
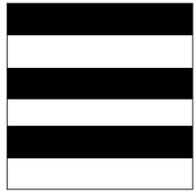


Figure B.5: Polar Coordinate Language Summary Charts & Examples;  $K = 10$

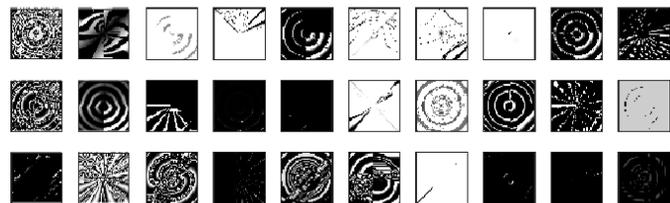
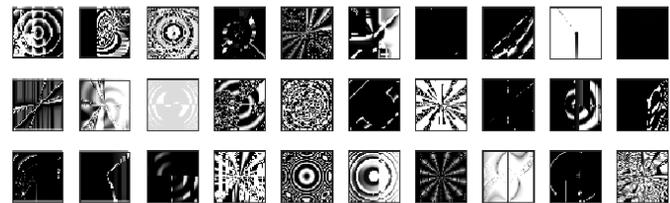
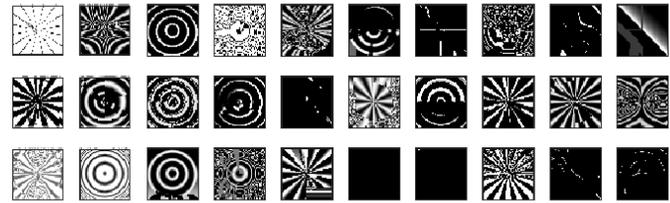
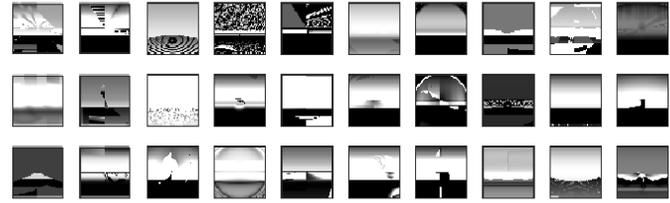
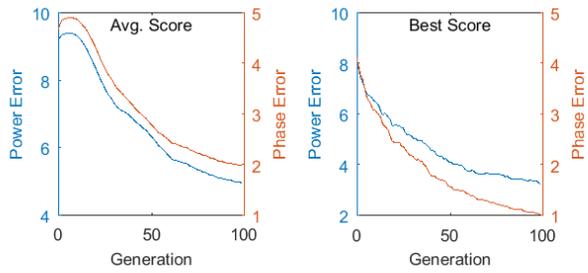
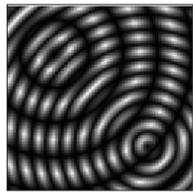
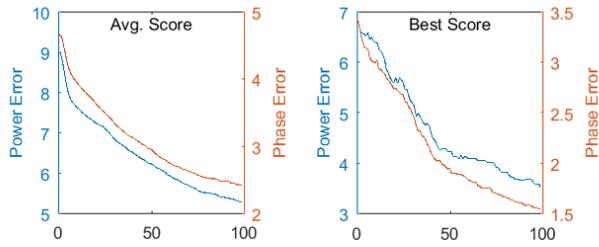
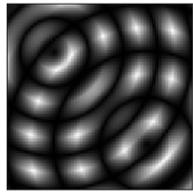
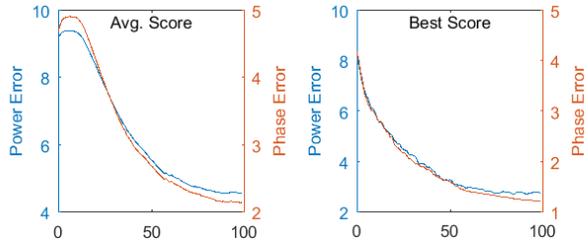
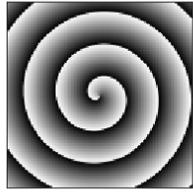
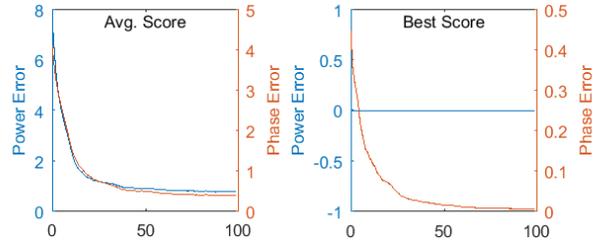


Figure B.6: Polar Coordinate Language Summary Charts & Examples;  $K = 25$

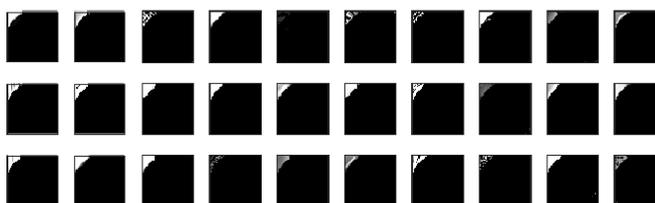
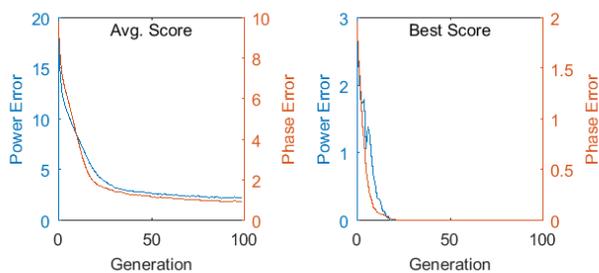
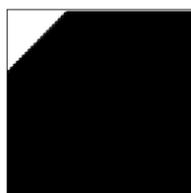
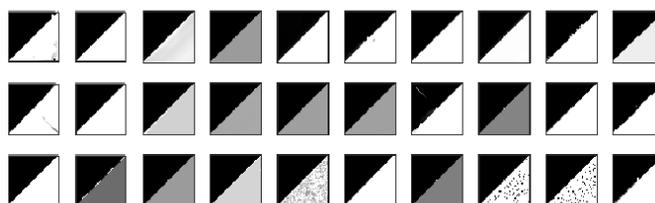
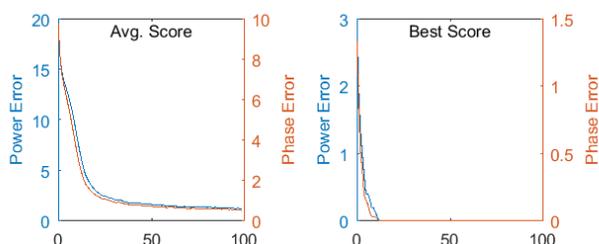
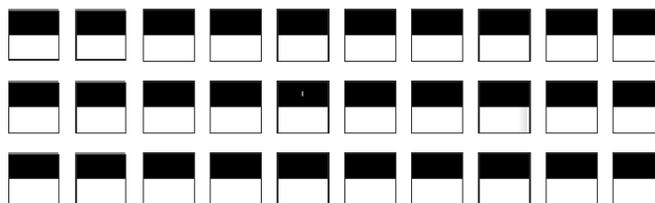
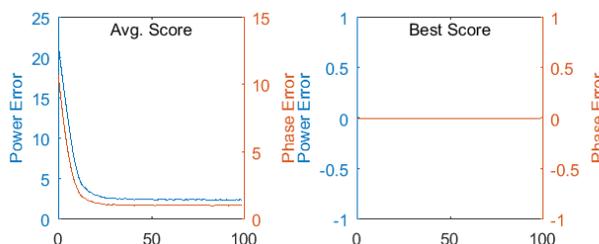
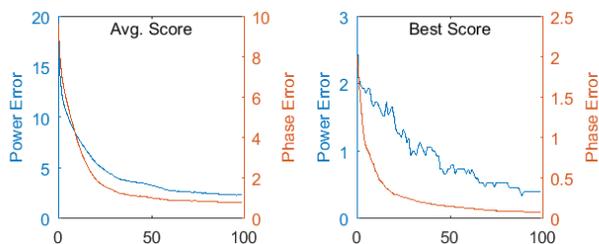
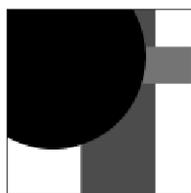


Figure B.6: Polar Coordinate Language Summary Charts & Examples;  $K = 25$

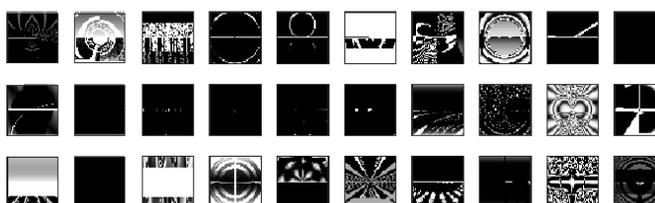
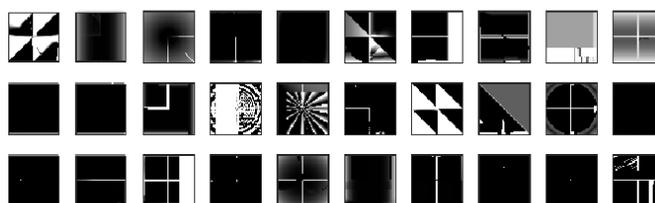
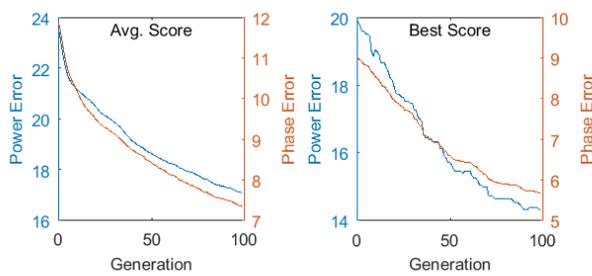
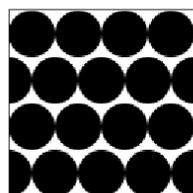
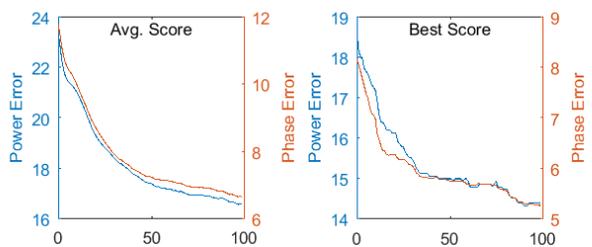
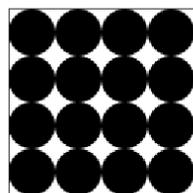
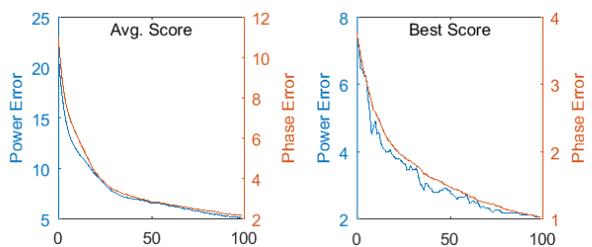
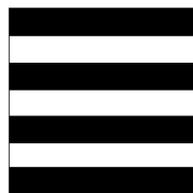
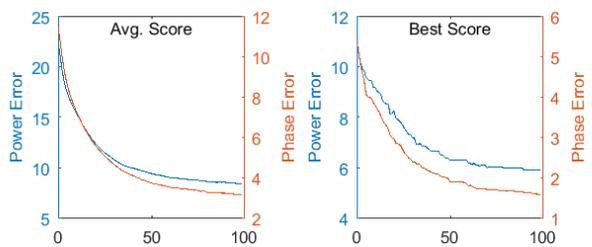


Figure B.6: Polar Coordinate Language Summary Charts & Examples;  $K = 25$

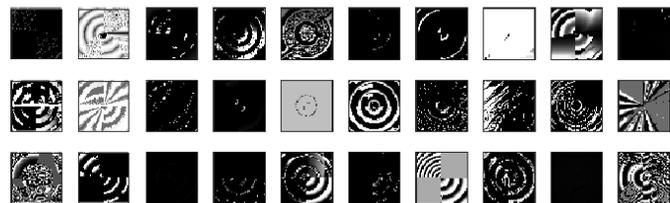
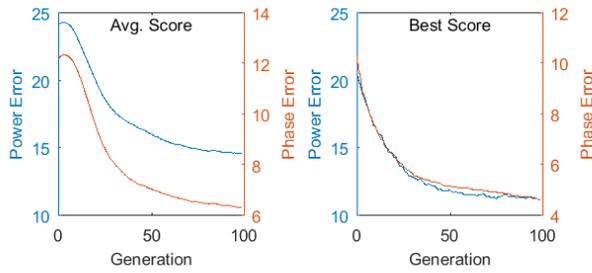
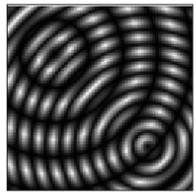
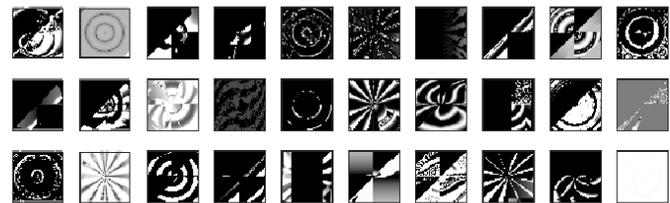
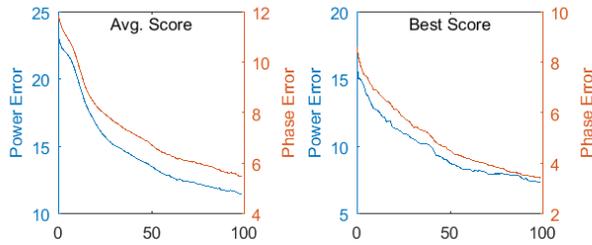
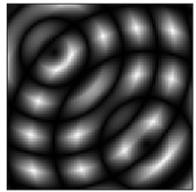
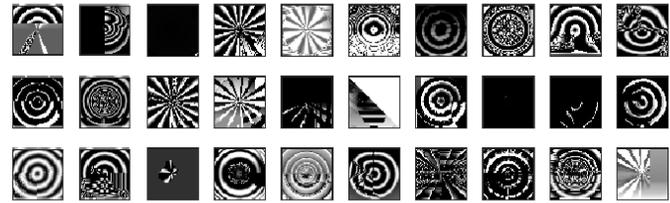
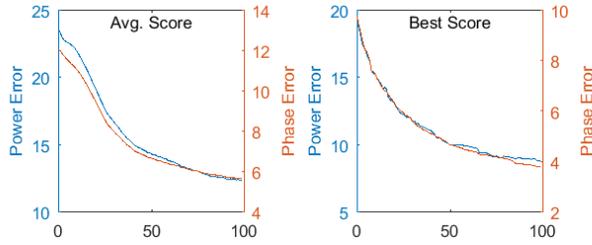
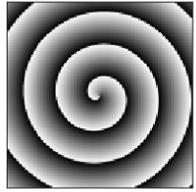
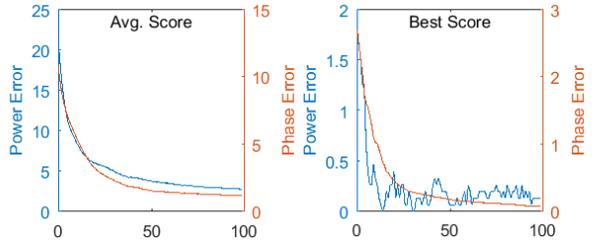


Figure B.7: Polar Coordinate Language Summary Charts & Examples;  $K = 50$

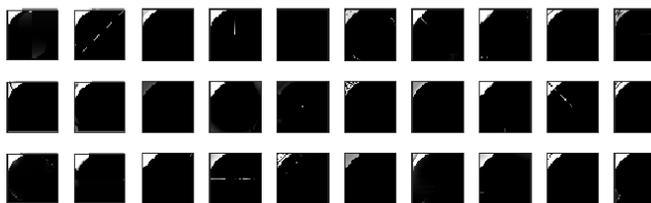
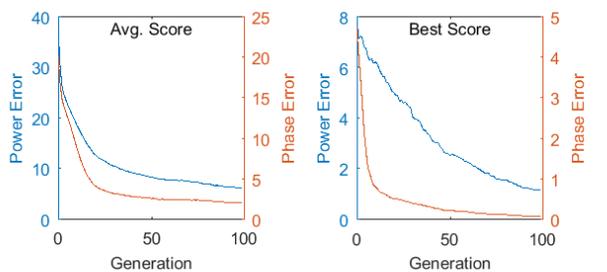
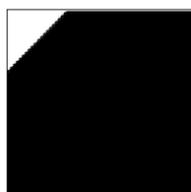
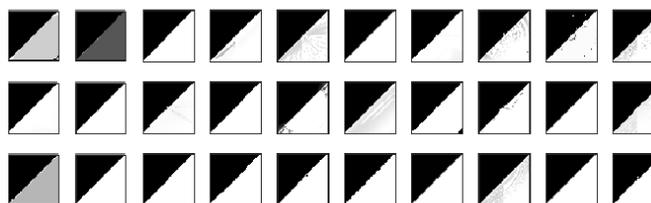
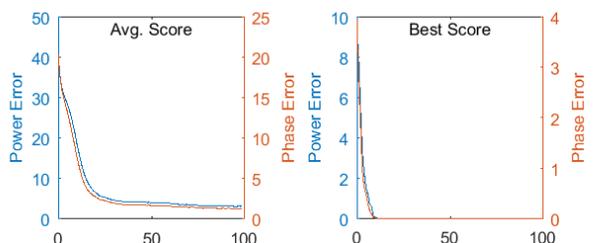
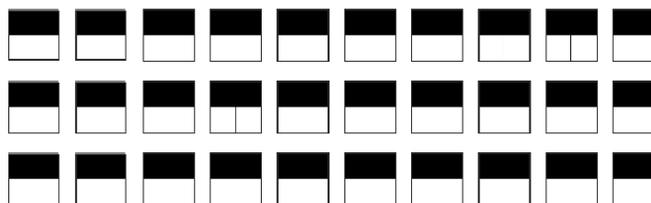
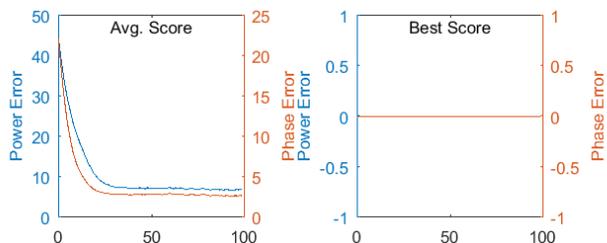
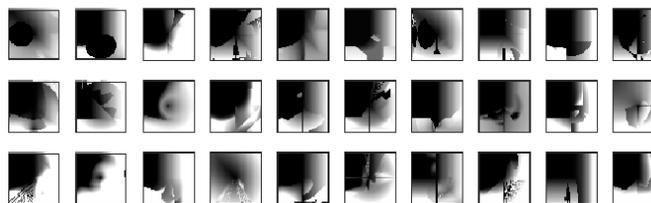
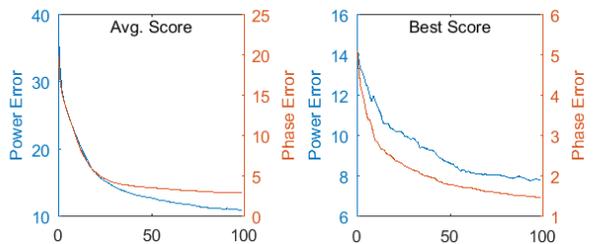
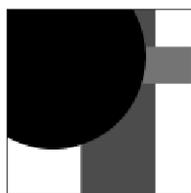


Figure B.7: Polar Coordinate Language Summary Charts & Examples;  $K = 50$

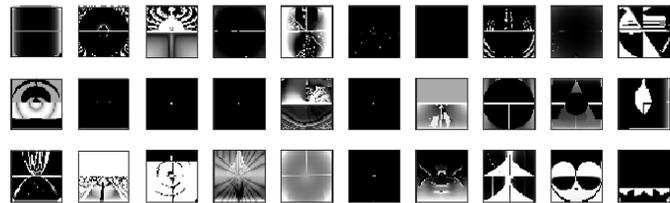
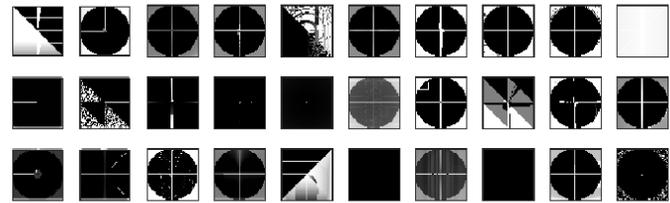
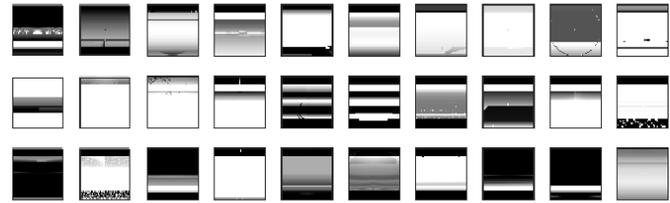
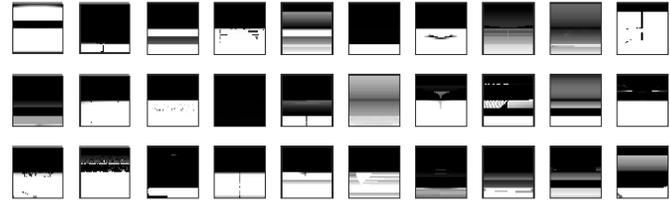
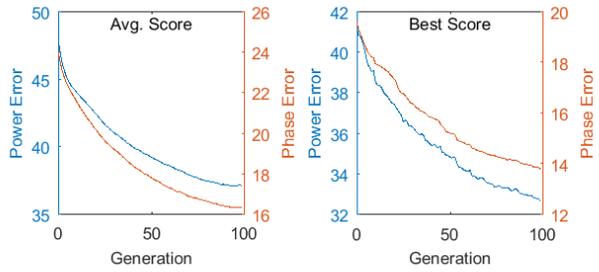
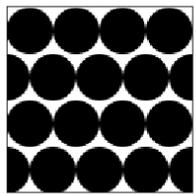
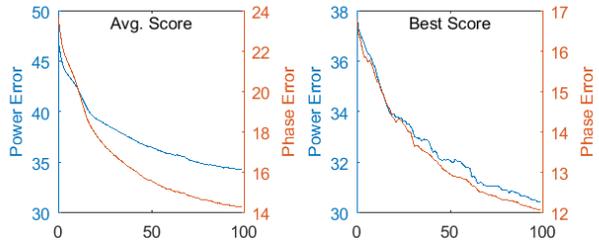
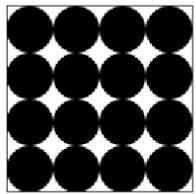
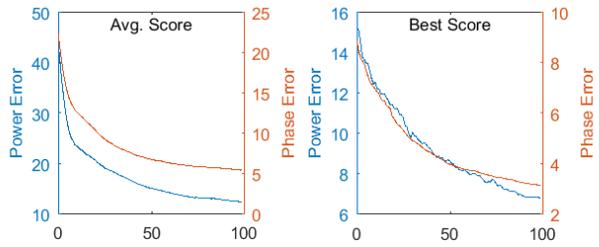
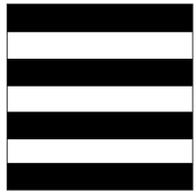
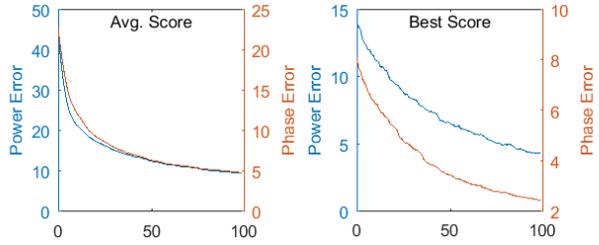
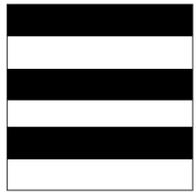


Figure B.7: Polar Coordinate Language Summary Charts & Examples;  $K = 50$

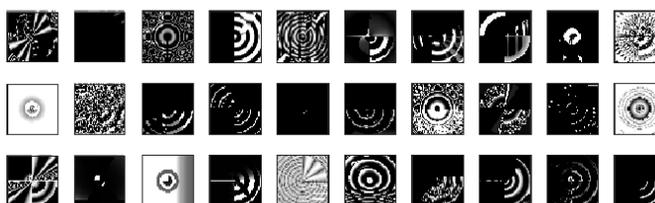
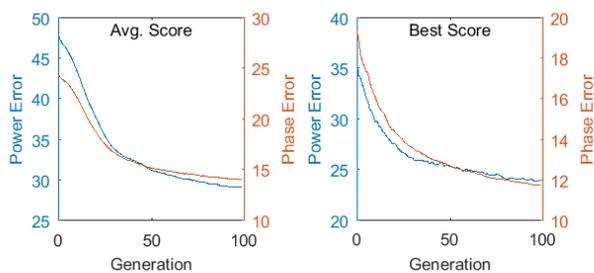
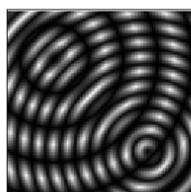
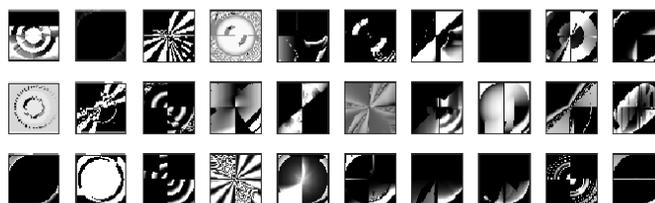
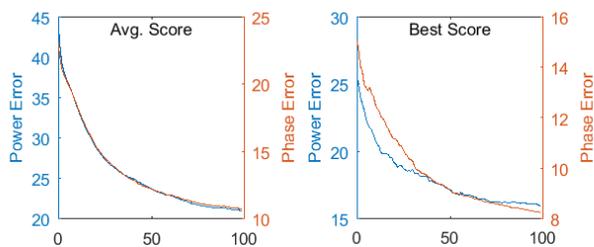
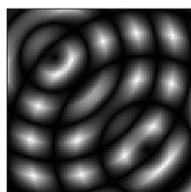
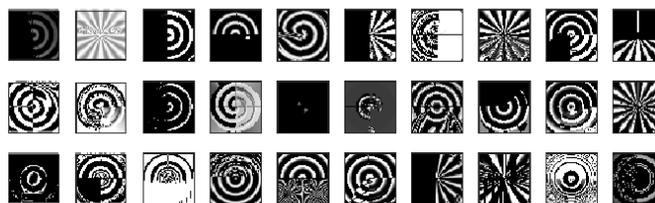
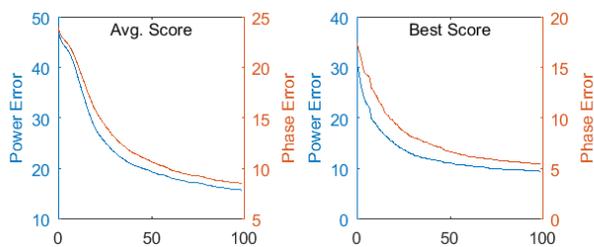
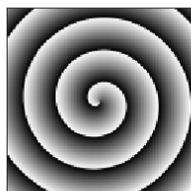
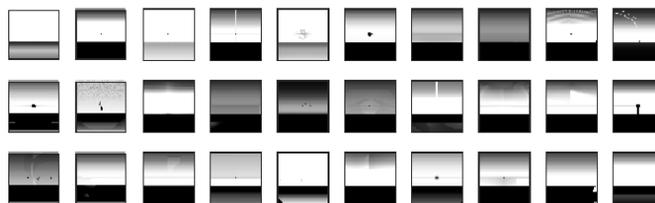
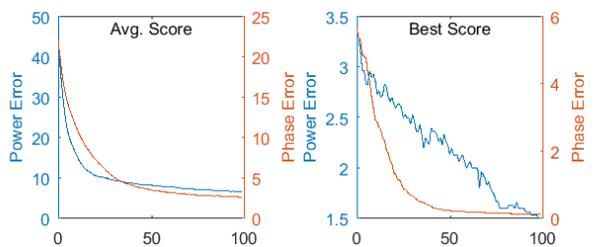


Table B.1: Polar Language Summary Table;  $K = 50$ 

| Target  | Agg.           | Power |       | Phase |       |       |
|---|----------------|-------|-------|-------|-------|-------|
|   |                | Mean  | Best  | Mean  | Best  |       |
|    | Composition_01 | Mean  | 10.98 | 7.80  | 2.94  | 1.48  |
|   | StdDev         | 1.15  | 1.37  | 0.42  | 0.29  |       |
|    | Composition_02 | Mean  | 6.83  | 0.00  | 2.65  | 0.00  |
|   | StdDev         | 1.05  | 0.00  | 0.48  | 0.00  |       |
|    | Composition_03 | Mean  | 3.27  | 0.00  | 1.35  | 0.00  |
|   | StdDev         | 1.19  | 0.00  | 0.55  | 0.00  |       |
|    | Composition_04 | Mean  | 6.32  | 1.20  | 2.12  | 0.08  |
|   | StdDev         | 1.52  | 1.40  | 0.63  | 0.12  |       |
|   | Composition_05 | Mean  | 9.54  | 4.33  | 4.84  | 2.43  |
|   | StdDev         | 2.16  | 2.48  | 0.92  | 1.02  |       |
|  | Composition_06 | Mean  | 12.45 | 6.80  | 5.51  | 3.15  |
|   | StdDev         | 2.92  | 3.21  | 1.13  | 0.99  |       |
|  | Composition_07 | Mean  | 34.36 | 30.47 | 14.31 | 12.09 |
|   | StdDev         | 3.26  | 4.44  | 1.16  | 1.46  |       |
|  | Composition_08 | Mean  | 37.15 | 32.67 | 16.35 | 13.81 |
|   | StdDev         | 3.07  | 4.74  | 1.74  | 2.51  |       |
|  | Composition_09 | Mean  | 6.66  | 1.50  | 2.65  | 0.13  |
|   | StdDev         | 1.33  | 0.86  | 0.77  | 0.10  |       |
|  | Composition_10 | Mean  | 15.79 | 9.50  | 8.51  | 5.51  |
|   | StdDev         | 4.82  | 4.66  | 2.89  | 3.04  |       |
|  | Composition_11 | Mean  | 21.07 | 15.97 | 10.79 | 8.26  |
|   | StdDev         | 2.14  | 2.85  | 1.30  | 1.55  |       |
|  | Composition_12 | Mean  | 29.10 | 23.93 | 14.04 | 11.76 |
|   | StdDev         | 3.79  | 4.70  | 1.60  | 1.82  |       |

Figure B.8: Circle, Gird, Offset Language Summary Charts & Examples;  $K = 50$

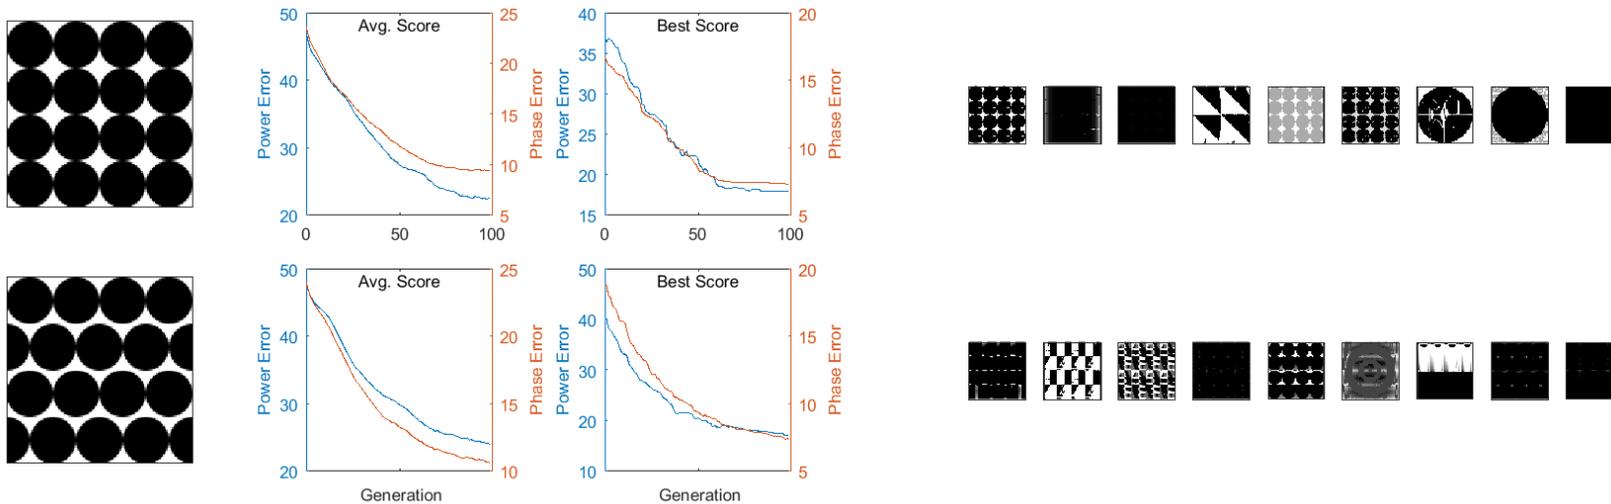


Figure B.9: Evolved Candidate Image With GP Tree - Grayscale

(a) Experiment P1;  $K = 50$ ; Noisy Language; Mondrian Target; Generation 60

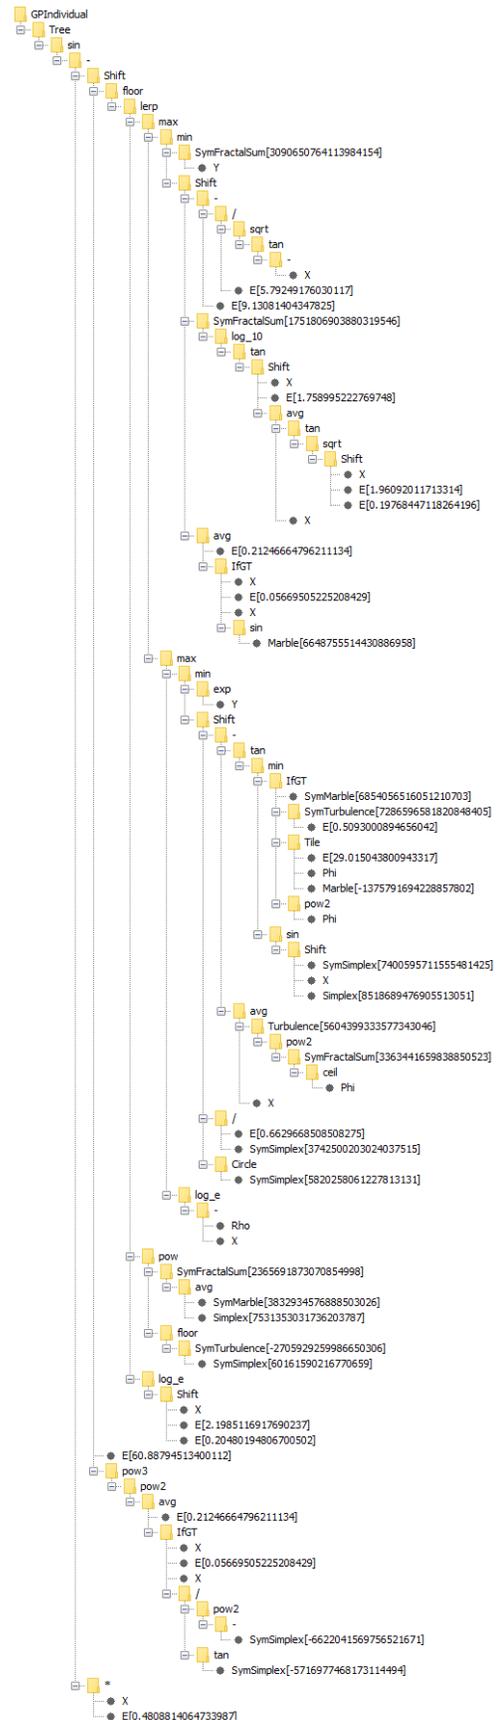
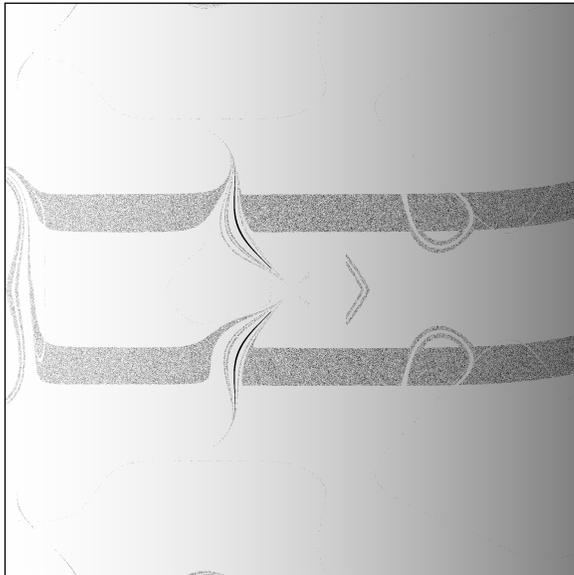
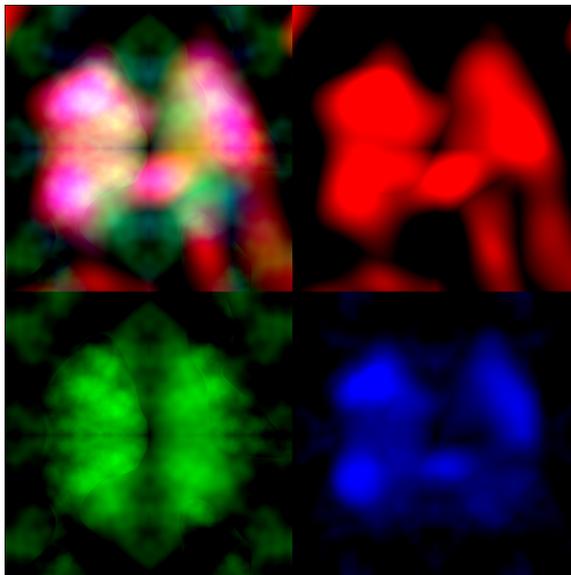


Figure B.10: Evolved Candidate Image With GP Tree - Colour

(a) Experiment C2;  $K = 50$ ; Noisy Language, No  $X$ ; Flower Target; Generation 100. Colour channels have been isolated and displayed alongside the evolved image.



# Appendix C

## User Survey & Data

A total of 36 complete sets of answers were acquired from the user survey. There were also 3 incomplete sets which have been removed, as exiting the browser had been an indicated method for users to withdraw participation consent.

We consider  $\mu = 0.5$ , comparing our mean and median against random selection chance.

Table C.1: User survey data and partial analysis

| Question | % Correct | $\Delta_{\mu_0}$ | Ordered | Signed Rank |
|----------|-----------|------------------|---------|-------------|
| 1        | 0.9722    | +0.4722          | +0.1389 | + 1         |
| 2        | 0.9444    | +0.4444          | +0.3056 | + 2         |
| 3        | 0.9722    | +0.4722          | +0.3611 | + 3         |
| 4        | 1.0000    | +0.5000          | +0.3889 | + 4         |
| 5        | 1.0000    | +0.5000          | +0.4167 | + 5         |
| 6        | 0.9444    | +0.4444          | +0.4444 | + 6 7       |
| 7        | 0.9722    | +0.4722          | +0.4444 | + 7 7       |
| 8        | 0.9444    | +0.4444          | +0.4444 | + 8 7       |
| 9        | 0.9167    | +0.4167          | +0.4722 | + 9 11      |
| 10       | 1.0000    | +0.5000          | +0.4722 | + 10 11     |
| 11       | 0.6389    | +0.1389          | +0.4722 | + 11 11     |
| 12       | 0.9722    | +0.4722          | +0.4722 | + 12 11     |
| 13       | 1.0000    | +0.5000          | +0.4722 | + 13 11     |
| 14       | 0.8611    | +0.3611          | +0.5000 | + 14 17     |
| 15       | 1.0000    | +0.5000          | +0.5000 | + 15 17     |
| 16       | 0.8889    | +0.3889          | +0.5000 | + 16 17     |
| 17       | 0.9722    | +0.4722          | +0.5000 | + 17 17     |
| 18       | 1.0000    | +0.5000          | +0.5000 | + 18 17     |
| 19       | 1.0000    | +0.5000          | +0.5000 | + 19 17     |
| 20       | 0.8056    | +0.3056          | +0.5000 | + 20 17     |

## Sign Test

The median for the null hypothesis should reflect a random selection. A reasonable confidence interval will be used.

$$\begin{aligned}\mu_0 &= 0.5 \\ \alpha &= 0.005\end{aligned}$$

We're comparing against 20 samples, none of which were excluded for matching  $\mu_0$ .

$$n = 20 - 0 = 20$$

For each sample, we obtain the difference from  $\mu_0$

We obtain a count of the positively signed differences, and a count of the negatively signed differences.

$$\begin{aligned}n^+ &= 20 \\ n^- &= 0\end{aligned}$$

The critical value to compare against for a double-tailed test would be the greater of the two counts. For our right-tailed test, we will use  $n^+$ .

$$\begin{aligned}T &= 20 \\ T_{n;\alpha} &= T_{20;0.005} = 14\end{aligned}$$

Since  $T > T_{n;\alpha}$ , we can reject the null hypothesis with a  $p \leq 0.005$  confidence level. Our observations are significant.

## Wilcoxon Signed-Rank Test

The median for the null hypothesis should reflect a random selection. A reasonable confidence interval will be used.

$$\begin{aligned}\mu_0 &= 0.5 \\ \alpha &= 0.005\end{aligned}$$

We're comparing against 20 samples, none of which were excluded for matching  $\mu_0$ .

$$n = 20 - 0 = 20$$

For each sample, we obtain the difference from  $\mu_0$ , and arrange them in ascending order by absolute value. Rankings are provided by this order, where groups of identical absolute values obtain the average rank of that group.

Determine the sum-of-ranks for positively signed differences, and the sum-of-ranks for negatively signed differences.

$$\begin{aligned}\Sigma R^+ &= 210 \\ \Sigma R^- &= 0\end{aligned}$$

The critical value to compare against for a double-tailed test would be the lesser of the two sum-of-rank values. For our right-tailed test, we will use  $\Sigma R^-$ .

$$\begin{aligned}T &= 0 \\ T_{n;\alpha} &= T_{20;0.005} = 14\end{aligned}$$

Since  $T < T_{n;\alpha}$ , we can reject the null hypothesis with a  $p \leq 0.005$  confidence level. Our observations are significant.

Survey: Evolving Textures Using 2D Power Spectra

INTRO

## Survey Evolving Textures Using 2D Power Spectra

### Principal Investigator, Faculty Supervisor

Dr. Brian Ross  
Department of Computer Science  
Brock University  
(905) 688-5550 Ext. 4284  
[bross@brocku.ca](mailto:bross@brocku.ca)

### Student Investigator

Michael Gircoys  
Department of Computer Science  
Brock University  
[mg12vp@brocku.ca](mailto:mg12vp@brocku.ca)

### Invitation

You are invited to participate in a study that involves research. The research will involve your participation in an online survey.

My ongoing MSc thesis research focuses on artificial intelligence software that generates abstract computer images. The computer software uses mathematical measurements to help generate new images that are visually similar to other existing images. The purpose of this online survey is to obtain some human opinions on the quality of the images generated by this system. In particular, we wish to determine how strong the similarity is between our computer-generated images and the existing images from a human's point of view. Our goal is to statistically determine whether the mathematical measures used by the software are useful in matching with human's opinions about the images.

### What's Involved

In this online survey, you will be presented with an existing art work, and two computer generated images. You will be asked to select the computer-generated image that you feel has a stronger visual similarity to the art work. There will be 20 example sets of images that you will be asked to evaluate, and the expected duration will be no more than 5 minutes in total. The survey will be conducted as an online questionnaire.

### Potential Benefits and Risks

This study will benefit the scientific community by providing statistical support to our mathematical models, correlating attributes in generated images with the abilities of human perception. It is hoped that these results can improve tools related to computational intelligence and machine vision.

There are no known or anticipated risks associated with participation in this study. The anonymous nature of the survey means that we cannot confirm or deny your participation in the survey.

### Confidentiality

Participants do not have to include any personal identification, such as name or email address. All recorded data will be anonymous. No information will be collected that will link your responses to your identity. All data collected during this study will be stored as part of my research and will be added to the appendix of my thesis. The digital data will be stored electronically in an encrypted archive on a local hard disk, and will be available to me and my supervisor. Although digital data may eventually be destroyed, completed survey data will reside in the thesis appendix for perpetuity. Note that incomplete surveys may not be used.

### Voluntary Participation

Participation in this study is voluntary. If you wish, you may decline to answer any questions or participate in any or all components of the study. Incomplete surveys will not be included in the research analysis. Further, you may decide to withdraw from this study at any time during the survey, and may do so without any penalty (e.g. you may close your browser at any time). However, it will not be possible to withdraw from participation after the survey has been completed, since it will not be possible to identify and remove a specific participant's responses from the collected data.

### Publication of Results

Participants can read about the results of this research in my completed thesis upon submission (anticipated in 2017/18), via the library's online collection. Results of this study may also be published in professional journals and presented at conferences. Data from this survey may be used by other researchers to validate results and figures within the thesis. Additional information about this study, an overview of results, as well as this thesis (upon completion), will be available at <http://www.cosc.brocku.ca/~bross/>.

### Contact Information and Ethics Clearance

If you have any questions about this study or require further information, please contact [Michael Gircoys](mailto:Michael.Gircoys@brocku.ca) and [Dr. Brian Ross](mailto:Dr.Brian.Ross@brocku.ca) using the contact information provided above. This study has been reviewed and received ethics clearance through the Research Ethics Board at Brock University. If you have any comments or concerns about your rights as a research participant, please contact the Research Ethics Office at (905) 688-5550 Ext. 3035, [reb@brocku.ca](mailto:reb@brocku.ca).

This research is partially supported by NSERC Discovery Grant 138467. This study has been reviewed and received ethics clearance through Brock University's Research Ethics Board (file #16-267).

Thank you for your assistance in this project.  
Please keep a copy of this form for your records.

### Consent Form

By clicking the web button below, I indicate that:

- I have read and understood the above information.
- I am 17 years of age or older.
- I agree to participate in this study. I have made this decision based on the information I have read in the Information-Consent Letter.
- I have had the opportunity to receive any additional details I wanted about the study and understand that I may ask questions in the future.
- I understand that I may withdraw this consent at any time during the survey.
- I should keep a copy of this consent form (web page) for my records.

There is no time limit to any question and no wrong answers, however, questions cannot be revisited.

Next

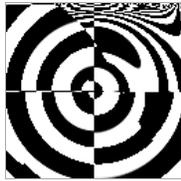
Survey: Evolving Textures Using 2D Power Spectra

1 / 20

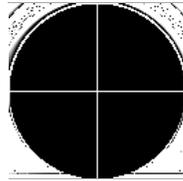


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

2 / 20

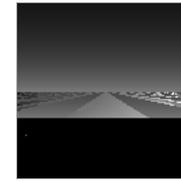


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

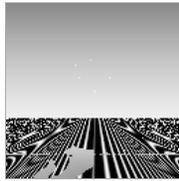
Survey: Evolving Textures Using 2D Power Spectra

3 / 20

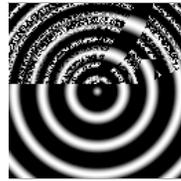


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

4 / 20

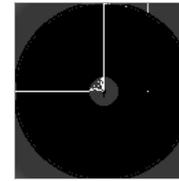


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

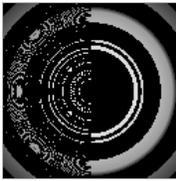
Survey: Evolving Textures Using 2D Power Spectra

5 / 20

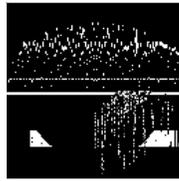


Target

Please select the image below which you feel is most similar to the target above.



Choice A

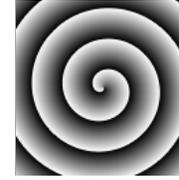


Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

6 / 20

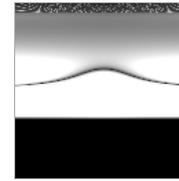


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

7 / 20

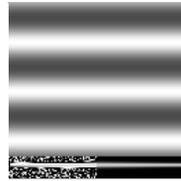


Target

Please select the image below which you feel is most similar to the target above.



Choice A

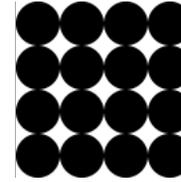


Choice B

Next

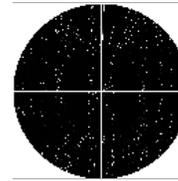
Survey: Evolving Textures Using 2D Power Spectra

8 / 20



Target

Please select the image below which you feel is most similar to the target above.



Choice A

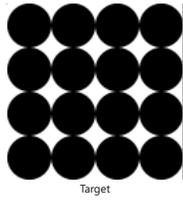


Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

9 / 20

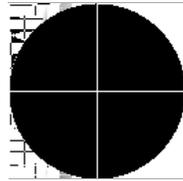


Target

Please select the image below which you feel is most similar to the target above.



Choice A

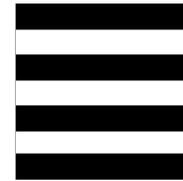


Choice B

Next

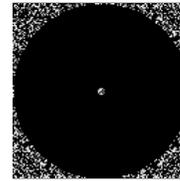
Survey: Evolving Textures Using 2D Power Spectra

10 / 20



Target

Please select the image below which you feel is most similar to the target above.



Choice A

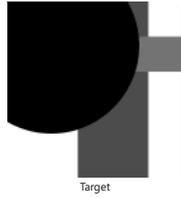


Choice B

Next

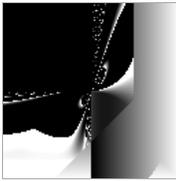
Survey: Evolving Textures Using 2D Power Spectra

11 / 20

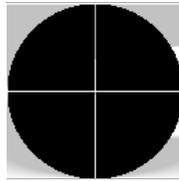


Target

Please select the image below which you feel is most similar to the target above.



Choice A

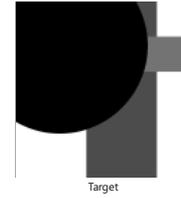


Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

12 / 20

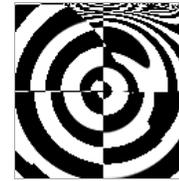


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

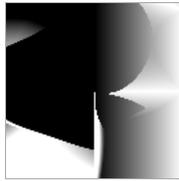
Survey: Evolving Textures Using 2D Power Spectra

13 / 20



Target

Please select the image below which you feel is most similar to the target above.



Choice A

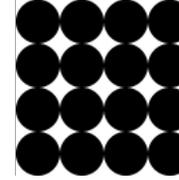


Choice B

Next

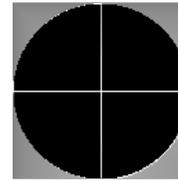
Survey: Evolving Textures Using 2D Power Spectra

14 / 20



Target

Please select the image below which you feel is most similar to the target above.



Choice A

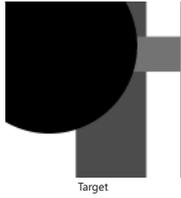


Choice B

Next

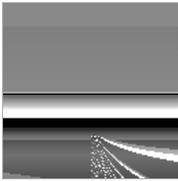
Survey: Evolving Textures Using 2D Power Spectra

15 / 20

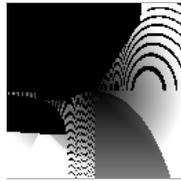


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

16 / 20

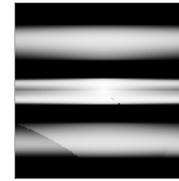


Target

Please select the image below which you feel is most similar to the target above.



Choice A



Choice B

Next

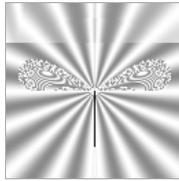
Survey: Evolving Textures Using 2D Power Spectra

17 / 20

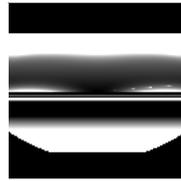


Target

Please select the image below which you feel is most similar to the target above.



Choice A

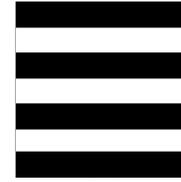


Choice B

Next

Survey: Evolving Textures Using 2D Power Spectra

18 / 20

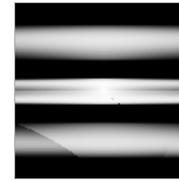


Target

Please select the image below which you feel is most similar to the target above.



Choice A

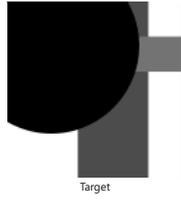


Choice B

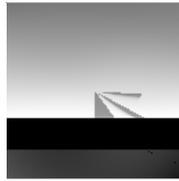
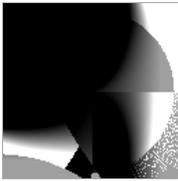
Next

Survey: Evolving Textures Using 2D Power Spectra

19 / 20



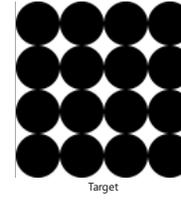
Please select the image below which you feel is most similar to the target above.



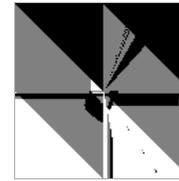
Next

Survey: Evolving Textures Using 2D Power Spectra

20 / 20



Please select the image below which you feel is most similar to the target above.



Next

Survey: Evolving Textures Using 2D Power Spectra

DONE

Thank you for your assistance in this project.

Additional information regarding this research project and other related works may be found at  
<http://www.cosc.brocku.ca/~bross/>

Table C.2: User survey raw data by question

| Session                    | Question | Correct | Session                    | Question | Correct | Session                    | Question | Correct |
|----------------------------|----------|---------|----------------------------|----------|---------|----------------------------|----------|---------|
| 0en70q9b8o6bn6bgfkhcf2i8h7 | 1        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7 | 3        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7 | 5        | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6 | 1        | 1       | 1u07p8omppjb2oqrtkadvdn2v6 | 3        | 1       | 1u07p8omppjb2oqrtkadvdn2v6 | 5        | 1       |
| 2mbikr6emk41b87kr46d68fh44 | 1        | 1       | 2mbikr6emk41b87kr46d68fh44 | 3        | 1       | 2mbikr6emk41b87kr46d68fh44 | 5        | 1       |
| 4pv1c0oi66fe06q2ck66l3o73  | 1        | 1       | 4pv1c0oi66fe06q2ck66l3o73  | 3        | 1       | 4pv1c0oi66fe06q2ck66l3o73  | 5        | 1       |
| 53tkvqq01r0jv30au4ikr2eb20 | 1        | 1       | 53tkvqq01r0jv30au4ikr2eb20 | 3        | 1       | 53tkvqq01r0jv30au4ikr2eb20 | 5        | 1       |
| 5f6po4ekr3bm57pik3md4tnub1 | 1        | 1       | 5f6po4ekr3bm57pik3md4tnub1 | 3        | 1       | 5f6po4ekr3bm57pik3md4tnub1 | 5        | 1       |
| 5n8u49lvg57r2pl76msom9s8v7 | 1        | 1       | 5n8u49lvg57r2pl76msom9s8v7 | 3        | 1       | 5n8u49lvg57r2pl76msom9s8v7 | 5        | 1       |
| 5rov1713rmqdn7mo4h715ln5n3 | 1        | 1       | 5rov1713rmqdn7mo4h715ln5n3 | 3        | 1       | 5rov1713rmqdn7mo4h715ln5n3 | 5        | 1       |
| 6rubc1hv5phadum9njf847an21 | 1        | 1       | 6rubc1hv5phadum9njf847an21 | 3        | 1       | 6rubc1hv5phadum9njf847an21 | 5        | 1       |
| 74bi9bd7o2ukkvfso9m18jou11 | 1        | 1       | 74bi9bd7o2ukkvfso9m18jou11 | 3        | 1       | 74bi9bd7o2ukkvfso9m18jou11 | 5        | 1       |
| 7rj4aa2crif7lshrson62ron41 | 1        | 1       | 7rj4aa2crif7lshrson62ron41 | 3        | 1       | 7rj4aa2crif7lshrson62ron41 | 5        | 1       |
| 8ejfs86cb24gludak26v5itop4 | 1        | 1       | 8ejfs86cb24gludak26v5itop4 | 3        | 1       | 8ejfs86cb24gludak26v5itop4 | 5        | 1       |
| 8ncopstrca2mqh9qmulkc79t4  | 1        | 1       | 8ncopstrca2mqh9qmulkc79t4  | 3        | 1       | 8ncopstrca2mqh9qmulkc79t4  | 5        | 1       |
| 8o1in7idjghsf9va07u06g1jk3 | 1        | 1       | 8o1in7idjghsf9va07u06g1jk3 | 3        | 1       | 8o1in7idjghsf9va07u06g1jk3 | 5        | 1       |
| 997f013mqp3p2h2qnp581cb76  | 1        | 1       | 997f013mqp3p2h2qnp581cb76  | 3        | 1       | 997f013mqp3p2h2qnp581cb76  | 5        | 1       |
| 9d1o1lp49nbnuilumoutrjtcq2 | 1        | 1       | 9d1o1lp49nbnuilumoutrjtcq2 | 3        | 1       | 9d1o1lp49nbnuilumoutrjtcq2 | 5        | 1       |
| 9m4evllj7tm1q9icls815o40m4 | 1        | 1       | 9m4evllj7tm1q9icls815o40m4 | 3        | 1       | 9m4evllj7tm1q9icls815o40m4 | 5        | 1       |
| biabpokjnk132lu2keb9nhmt5  | 1        | 1       | biabpokjnk132lu2keb9nhmt5  | 3        | 1       | biabpokjnk132lu2keb9nhmt5  | 5        | 1       |
| c7bmbkeriqjfr85sm3vcbth933 | 1        | 1       | c7bmbkeriqjfr85sm3vcbth933 | 3        | 1       | c7bmbkeriqjfr85sm3vcbth933 | 5        | 1       |
| csli5u3ms3rtg8qht0rfihd6q3 | 1        | 1       | csli5u3ms3rtg8qht0rfihd6q3 | 3        | 1       | csli5u3ms3rtg8qht0rfihd6q3 | 5        | 1       |
| e0s8q070do39qcm87s3mktbfj6 | 1        | 1       | e0s8q070do39qcm87s3mktbfj6 | 3        | 1       | e0s8q070do39qcm87s3mktbfj6 | 5        | 1       |
| euc6ho5uqo8vputqo7u0jnihc1 | 1        | 1       | euc6ho5uqo8vputqo7u0jnihc1 | 3        | 1       | euc6ho5uqo8vputqo7u0jnihc1 | 5        | 1       |
| frj8s3fd1po84uje7sftrh85m5 | 1        | 1       | frj8s3fd1po84uje7sftrh85m5 | 3        | 1       | frj8s3fd1po84uje7sftrh85m5 | 5        | 1       |
| gj0e3kmm5tuaasosvndri1663  | 1        | 1       | gj0e3kmm5tuaasosvndri1663  | 3        | 1       | gj0e3kmm5tuaasosvndri1663  | 5        | 1       |
| i97epapdrhnh3sisupuju7orb0 | 1        | 1       | i97epapdrhnh3sisupuju7orb0 | 3        | 1       | i97epapdrhnh3sisupuju7orb0 | 5        | 1       |
| ij38i1664l1bmd2lrorlts4jn2 | 1        | 1       | ij38i1664l1bmd2lrorlts4jn2 | 3        | 1       | ij38i1664l1bmd2lrorlts4jn2 | 5        | 1       |
| msdnp47br71gdv11iuidgnrt43 | 1        | 1       | msdnp47br71gdv11iuidgnrt43 | 3        | 1       | msdnp47br71gdv11iuidgnrt43 | 5        | 1       |
| nfnkmvflpcukq18pvd5ds4r22  | 1        | 1       | nfnkmvflpcukq18pvd5ds4r22  | 3        | 1       | nfnkmvflpcukq18pvd5ds4r22  | 5        | 1       |
| o4ceqglln1nphd7c6kb4b93923 | 1        | 1       | o4ceqglln1nphd7c6kb4b93923 | 3        | 0       | o4ceqglln1nphd7c6kb4b93923 | 5        | 1       |
| oep8758jc2ae7983kml3qqupf2 | 1        | 1       | oep8758jc2ae7983kml3qqupf2 | 3        | 1       | oep8758jc2ae7983kml3qqupf2 | 5        | 1       |
| p3beelqe1gbd8tk2qic2hbfj5  | 1        | 1       | p3beelqe1gbd8tk2qic2hbfj5  | 3        | 1       | p3beelqe1gbd8tk2qic2hbfj5  | 5        | 1       |
| qt4orejugu7hqbd0m8jr7j9d56 | 1        | 1       | qt4orejugu7hqbd0m8jr7j9d56 | 3        | 1       | qt4orejugu7hqbd0m8jr7j9d56 | 5        | 1       |
| rde6dsetg9tn721f8ef11giug5 | 1        | 0       | rde6dsetg9tn721f8ef11giug5 | 3        | 1       | rde6dsetg9tn721f8ef11giug5 | 5        | 1       |
| risfv5f45sjk2k0hsjhb0cnp2  | 1        | 1       | risfv5f45sjk2k0hsjhb0cnp2  | 3        | 1       | risfv5f45sjk2k0hsjhb0cnp2  | 5        | 1       |
| umbkj3outn2q0u4og1inh0d86  | 1        | 1       | umbkj3outn2q0u4og1inh0d86  | 3        | 1       | umbkj3outn2q0u4og1inh0d86  | 5        | 1       |
| vulccmpfd9kigtmm76qvekuk11 | 1        | 1       | vulccmpfd9kigtmm76qvekuk11 | 3        | 1       | vulccmpfd9kigtmm76qvekuk11 | 5        | 1       |
| 0en70q9b8o6bn6bgfkhcf2i8h7 | 2        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7 | 4        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7 | 6        | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6 | 2        | 0       | 1u07p8omppjb2oqrtkadvdn2v6 | 4        | 1       | 1u07p8omppjb2oqrtkadvdn2v6 | 6        | 1       |
| 2mbikr6emk41b87kr46d68fh44 | 2        | 1       | 2mbikr6emk41b87kr46d68fh44 | 4        | 1       | 2mbikr6emk41b87kr46d68fh44 | 6        | 1       |
| 4pv1c0oi66fe06q2ck66l3o73  | 2        | 1       | 4pv1c0oi66fe06q2ck66l3o73  | 4        | 1       | 4pv1c0oi66fe06q2ck66l3o73  | 6        | 1       |
| 53tkvqq01r0jv30au4ikr2eb20 | 2        | 1       | 53tkvqq01r0jv30au4ikr2eb20 | 4        | 1       | 53tkvqq01r0jv30au4ikr2eb20 | 6        | 1       |
| 5f6po4ekr3bm57pik3md4tnub1 | 2        | 1       | 5f6po4ekr3bm57pik3md4tnub1 | 4        | 1       | 5f6po4ekr3bm57pik3md4tnub1 | 6        | 1       |
| 5n8u49lvg57r2pl76msom9s8v7 | 2        | 1       | 5n8u49lvg57r2pl76msom9s8v7 | 4        | 1       | 5n8u49lvg57r2pl76msom9s8v7 | 6        | 1       |
| 5rov1713rmqdn7mo4h715ln5n3 | 2        | 1       | 5rov1713rmqdn7mo4h715ln5n3 | 4        | 1       | 5rov1713rmqdn7mo4h715ln5n3 | 6        | 1       |
| 6rubc1hv5phadum9njf847an21 | 2        | 1       | 6rubc1hv5phadum9njf847an21 | 4        | 1       | 6rubc1hv5phadum9njf847an21 | 6        | 1       |
| 74bi9bd7o2ukkvfso9m18jou11 | 2        | 1       | 74bi9bd7o2ukkvfso9m18jou11 | 4        | 1       | 74bi9bd7o2ukkvfso9m18jou11 | 6        | 1       |
| 7rj4aa2crif7lshrson62ron41 | 2        | 1       | 7rj4aa2crif7lshrson62ron41 | 4        | 1       | 7rj4aa2crif7lshrson62ron41 | 6        | 1       |
| 8ejfs86cb24gludak26v5itop4 | 2        | 1       | 8ejfs86cb24gludak26v5itop4 | 4        | 1       | 8ejfs86cb24gludak26v5itop4 | 6        | 1       |
| 8ncopstrca2mqh9qmulkc79t4  | 2        | 1       | 8ncopstrca2mqh9qmulkc79t4  | 4        | 1       | 8ncopstrca2mqh9qmulkc79t4  | 6        | 1       |
| 8o1in7idjghsf9va07u06g1jk3 | 2        | 1       | 8o1in7idjghsf9va07u06g1jk3 | 4        | 1       | 8o1in7idjghsf9va07u06g1jk3 | 6        | 1       |
| 997f013mqp3p2h2qnp581cb76  | 2        | 1       | 997f013mqp3p2h2qnp581cb76  | 4        | 1       | 997f013mqp3p2h2qnp581cb76  | 6        | 1       |
| 9d1o1lp49nbnuilumoutrjtcq2 | 2        | 1       | 9d1o1lp49nbnuilumoutrjtcq2 | 4        | 1       | 9d1o1lp49nbnuilumoutrjtcq2 | 6        | 0       |
| 9m4evllj7tm1q9icls815o40m4 | 2        | 1       | 9m4evllj7tm1q9icls815o40m4 | 4        | 1       | 9m4evllj7tm1q9icls815o40m4 | 6        | 1       |
| biabpokjnk132lu2keb9nhmt5  | 2        | 0       | biabpokjnk132lu2keb9nhmt5  | 4        | 1       | biabpokjnk132lu2keb9nhmt5  | 6        | 1       |
| c7bmbkeriqjfr85sm3vcbth933 | 2        | 1       | c7bmbkeriqjfr85sm3vcbth933 | 4        | 1       | c7bmbkeriqjfr85sm3vcbth933 | 6        | 1       |
| csli5u3ms3rtg8qht0rfihd6q3 | 2        | 1       | csli5u3ms3rtg8qht0rfihd6q3 | 4        | 1       | csli5u3ms3rtg8qht0rfihd6q3 | 6        | 1       |
| e0s8q070do39qcm87s3mktbfj6 | 2        | 1       | e0s8q070do39qcm87s3mktbfj6 | 4        | 1       | e0s8q070do39qcm87s3mktbfj6 | 6        | 1       |
| euc6ho5uqo8vputqo7u0jnihc1 | 2        | 1       | euc6ho5uqo8vputqo7u0jnihc1 | 4        | 1       | euc6ho5uqo8vputqo7u0jnihc1 | 6        | 1       |
| frj8s3fd1po84uje7sftrh85m5 | 2        | 1       | frj8s3fd1po84uje7sftrh85m5 | 4        | 1       | frj8s3fd1po84uje7sftrh85m5 | 6        | 1       |
| gj0e3kmm5tuaasosvndri1663  | 2        | 1       | gj0e3kmm5tuaasosvndri1663  | 4        | 1       | gj0e3kmm5tuaasosvndri1663  | 6        | 1       |
| i97epapdrhnh3sisupuju7orb0 | 2        | 1       | i97epapdrhnh3sisupuju7orb0 | 4        | 1       | i97epapdrhnh3sisupuju7orb0 | 6        | 1       |
| ij38i1664l1bmd2lrorlts4jn2 | 2        | 1       | ij38i1664l1bmd2lrorlts4jn2 | 4        | 1       | ij38i1664l1bmd2lrorlts4jn2 | 6        | 0       |
| msdnp47br71gdv11iuidgnrt43 | 2        | 1       | msdnp47br71gdv11iuidgnrt43 | 4        | 1       | msdnp47br71gdv11iuidgnrt43 | 6        | 1       |
| nfnkmvflpcukq18pvd5ds4r22  | 2        | 1       | nfnkmvflpcukq18pvd5ds4r22  | 4        | 1       | nfnkmvflpcukq18pvd5ds4r22  | 6        | 1       |
| o4ceqglln1nphd7c6kb4b93923 | 2        | 1       | o4ceqglln1nphd7c6kb4b93923 | 4        | 1       | o4ceqglln1nphd7c6kb4b93923 | 6        | 1       |
| oep8758jc2ae7983kml3qqupf2 | 2        | 1       | oep8758jc2ae7983kml3qqupf2 | 4        | 1       | oep8758jc2ae7983kml3qqupf2 | 6        | 1       |
| p3beelqe1gbd8tk2qic2hbfj5  | 2        | 1       | p3beelqe1gbd8tk2qic2hbfj5  | 4        | 1       | p3beelqe1gbd8tk2qic2hbfj5  | 6        | 1       |
| qt4orejugu7hqbd0m8jr7j9d56 | 2        | 1       | qt4orejugu7hqbd0m8jr7j9d56 | 4        | 1       | qt4orejugu7hqbd0m8jr7j9d56 | 6        | 1       |
| rde6dsetg9tn721f8ef11giug5 | 2        | 1       | rde6dsetg9tn721f8ef11giug5 | 4        | 1       | rde6dsetg9tn721f8ef11giug5 | 6        | 1       |
| risfv5f45sjk2k0hsjhb0cnp2  | 2        | 1       | risfv5f45sjk2k0hsjhb0cnp2  | 4        | 1       | risfv5f45sjk2k0hsjhb0cnp2  | 6        | 1       |
| umbkj3outn2q0u4og1inh0d86  | 2        | 1       | umbkj3outn2q0u4og1inh0d86  | 4        | 1       | umbkj3outn2q0u4og1inh0d86  | 6        | 1       |
| vulccmpfd9kigtmm76qvekuk11 | 2        | 1       | vulccmpfd9kigtmm76qvekuk11 | 4        | 1       | vulccmpfd9kigtmm76qvekuk11 | 6        | 1       |

| Session                     | Question | Correct | Session                     | Question | Correct | Session                     | Question | Correct |
|-----------------------------|----------|---------|-----------------------------|----------|---------|-----------------------------|----------|---------|
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 7        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 9        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 11       | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 7        | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 9        | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 11       | 1       |
| 2mbikr6emk41b87kr46d68fh44  | 7        | 1       | 2mbikr6emk41b87kr46d68fh44  | 9        | 0       | 2mbikr6emk41b87kr46d68fh44  | 11       | 0       |
| 4pv1c0oi66fe06q2ck66l3o73   | 7        | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 9        | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 11       | 0       |
| 53tkvqq01r0jv30au4ikr2eb20  | 7        | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 9        | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 11       | 1       |
| 5f6po4ekr3bm57pik3md4tnub1  | 7        | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 9        | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 11       | 1       |
| 5n8u49lvgs7r2pl76msom9s8v7  | 7        | 1       | 5n8u49lvgs7r2pl76msom9s8v7  | 9        | 1       | 5n8u49lvgs7r2pl76msom9s8v7  | 11       | 1       |
| 5rov1713rmqdn7mo4h715ln5n3  | 7        | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 9        | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 11       | 0       |
| 6rubc1hv5phadam9njf847an21  | 7        | 1       | 6rubc1hv5phadam9njf847an21  | 9        | 1       | 6rubc1hv5phadam9njf847an21  | 11       | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 7        | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 9        | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 11       | 1       |
| 7rj4aa2crif7lshrson62ron41  | 7        | 1       | 7rj4aa2crif7lshrson62ron41  | 9        | 1       | 7rj4aa2crif7lshrson62ron41  | 11       | 0       |
| 8ejfs86cb24gludak26v5itop4  | 7        | 1       | 8ejfs86cb24gludak26v5itop4  | 9        | 1       | 8ejfs86cb24gludak26v5itop4  | 11       | 0       |
| 8ncopsrtca2mqqh9qmulkc79t4  | 7        | 1       | 8ncopsrtca2mqqh9qmulkc79t4  | 9        | 1       | 8ncopsrtca2mqqh9qmulkc79t4  | 11       | 1       |
| 8o1in7idjqhsf9va07uo6g1jk3  | 7        | 1       | 8o1in7idjqhsf9va07uo6g1jk3  | 9        | 1       | 8o1in7idjqhsf9va07uo6g1jk3  | 11       | 1       |
| 997f013mqp3p2h2qnpm581cb76  | 7        | 1       | 997f013mqp3p2h2qnpm581cb76  | 9        | 1       | 997f013mqp3p2h2qnpm581cb76  | 11       | 1       |
| 9d1o1lp49nbuilmoutrjtcq2    | 7        | 1       | 9d1o1lp49nbuilmoutrjtcq2    | 9        | 1       | 9d1o1lp49nbuilmoutrjtcq2    | 11       | 1       |
| 9m4evllj7tm1q9icls815o40m4  | 7        | 1       | 9m4evllj7tm1q9icls815o40m4  | 9        | 1       | 9m4evllj7tm1q9icls815o40m4  | 11       | 0       |
| biabpokjnk132lu2keb9nhmt5   | 7        | 1       | biabpokjnk132lu2keb9nhmt5   | 9        | 1       | biabpokjnk132lu2keb9nhmt5   | 11       | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 7        | 1       | c7bmbkeriqjfr85sm3vcbth933  | 9        | 1       | c7bmbkeriqjfr85sm3vcbth933  | 11       | 0       |
| csli5u3ms3rtg8qht0rfihd6q3  | 7        | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 9        | 0       | csli5u3ms3rtg8qht0rfihd6q3  | 11       | 0       |
| e0s8q070do39qcm87s3mktbfj6  | 7        | 1       | e0s8q070do39qcm87s3mktbfj6  | 9        | 1       | e0s8q070do39qcm87s3mktbfj6  | 11       | 0       |
| euc6ho5uqo8vputqo7u0jnihc1  | 7        | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 9        | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 11       | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 7        | 1       | frj8s3fd1po84uje7sftrh85m5  | 9        | 1       | frj8s3fd1po84uje7sftrh85m5  | 11       | 1       |
| gj0e3kmm5tuaasosvni dri1663 | 7        | 1       | gj0e3kmm5tuaasosvni dri1663 | 9        | 1       | gj0e3kmm5tuaasosvni dri1663 | 11       | 1       |
| i97epapdrhnh3sisupuju7orb0  | 7        | 1       | i97epapdrhnh3sisupuju7orb0  | 9        | 1       | i97epapdrhnh3sisupuju7orb0  | 11       | 1       |
| ij381i664l bmd2lrorlts4jn2  | 7        | 0       | ij381i664l bmd2lrorlts4jn2  | 9        | 1       | ij381i664l bmd2lrorlts4jn2  | 11       | 1       |
| msdnp47br71gdv11iuidgnrt43  | 7        | 1       | msdnp47br71gdv11iuidgnrt43  | 9        | 1       | msdnp47br71gdv11iuidgnrt43  | 11       | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 7        | 1       | nfnkmvflpcukq18pvdssds4r22  | 9        | 1       | nfnkmvflpcukq18pvdssds4r22  | 11       | 0       |
| o4ceqglln1nphd7c6kb4b93923  | 7        | 1       | o4ceqglln1nphd7c6kb4b93923  | 9        | 0       | o4ceqglln1nphd7c6kb4b93923  | 11       | 0       |
| oep8758jc2ae7983kml3qqupf2  | 7        | 1       | oep8758jc2ae7983kml3qqupf2  | 9        | 1       | oep8758jc2ae7983kml3qqupf2  | 11       | 1       |
| p3beelqe1gbd8tk2qic2hbfej5  | 7        | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 9        | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 11       | 1       |
| qt4orejugu7hqbd0m8j7r7j9d56 | 7        | 1       | qt4orejugu7hqbd0m8j7r7j9d56 | 9        | 1       | qt4orejugu7hqbd0m8j7r7j9d56 | 11       | 1       |
| rde6dsetg9tn721f8ef11giug5  | 7        | 1       | rde6dsetg9tn721f8ef11giug5  | 9        | 1       | rde6dsetg9tn721f8ef11giug5  | 11       | 1       |
| risfv5f45skj2k0hsjhb0cnp2   | 7        | 1       | risfv5f45skj2k0hsjhb0cnp2   | 9        | 1       | risfv5f45skj2k0hsjhb0cnp2   | 11       | 0       |
| umbkj3outn2q0u4ogi1nhm0d86  | 7        | 1       | umbkj3outn2q0u4ogi1nhm0d86  | 9        | 1       | umbkj3outn2q0u4ogi1nhm0d86  | 11       | 0       |
| vulccmpfd9kigtmm76qvekukl1  | 7        | 1       | vulccmpfd9kigtmm76qvekukl1  | 9        | 1       | vulccmpfd9kigtmm76qvekukl1  | 11       | 0       |
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 8        | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 10       | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 12       | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 8        | 0       | 1u07p8omppjb2oqrtkadvdn2v6  | 10       | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 12       | 1       |
| 2mbikr6emk41b87kr46d68fh44  | 8        | 1       | 2mbikr6emk41b87kr46d68fh44  | 10       | 1       | 2mbikr6emk41b87kr46d68fh44  | 12       | 1       |
| 4pv1c0oi66fe06q2ck66l3o73   | 8        | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 10       | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 12       | 1       |
| 53tkvqq01r0jv30au4ikr2eb20  | 8        | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 10       | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 12       | 1       |
| 5f6po4ekr3bm57pik3md4tnub1  | 8        | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 10       | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 12       | 1       |
| 5n8u49lvgs7r2pl76msom9s8v7  | 8        | 1       | 5n8u49lvgs7r2pl76msom9s8v7  | 10       | 1       | 5n8u49lvgs7r2pl76msom9s8v7  | 12       | 1       |
| 5rov1713rmqdn7mo4h715ln5n3  | 8        | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 10       | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 12       | 1       |
| 6rubc1hv5phadam9njf847an21  | 8        | 1       | 6rubc1hv5phadam9njf847an21  | 10       | 1       | 6rubc1hv5phadam9njf847an21  | 12       | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 8        | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 10       | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 12       | 1       |
| 7rj4aa2crif7lshrson62ron41  | 8        | 1       | 7rj4aa2crif7lshrson62ron41  | 10       | 1       | 7rj4aa2crif7lshrson62ron41  | 12       | 1       |
| 8ejfs86cb24gludak26v5itop4  | 8        | 1       | 8ejfs86cb24gludak26v5itop4  | 10       | 1       | 8ejfs86cb24gludak26v5itop4  | 12       | 1       |
| 8ncopsrtca2mqqh9qmulkc79t4  | 8        | 1       | 8ncopsrtca2mqqh9qmulkc79t4  | 10       | 1       | 8ncopsrtca2mqqh9qmulkc79t4  | 12       | 1       |
| 8o1in7idjqhsf9va07uo6g1jk3  | 8        | 1       | 8o1in7idjqhsf9va07uo6g1jk3  | 10       | 1       | 8o1in7idjqhsf9va07uo6g1jk3  | 12       | 1       |
| 997f013mqp3p2h2qnpm581cb76  | 8        | 1       | 997f013mqp3p2h2qnpm581cb76  | 10       | 1       | 997f013mqp3p2h2qnpm581cb76  | 12       | 1       |
| 9d1o1lp49nbuilmoutrjtcq2    | 8        | 1       | 9d1o1lp49nbuilmoutrjtcq2    | 10       | 1       | 9d1o1lp49nbuilmoutrjtcq2    | 12       | 1       |
| 9m4evllj7tm1q9icls815o40m4  | 8        | 1       | 9m4evllj7tm1q9icls815o40m4  | 10       | 1       | 9m4evllj7tm1q9icls815o40m4  | 12       | 1       |
| biabpokjnk132lu2keb9nhmt5   | 8        | 1       | biabpokjnk132lu2keb9nhmt5   | 10       | 1       | biabpokjnk132lu2keb9nhmt5   | 12       | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 8        | 1       | c7bmbkeriqjfr85sm3vcbth933  | 10       | 1       | c7bmbkeriqjfr85sm3vcbth933  | 12       | 0       |
| csli5u3ms3rtg8qht0rfihd6q3  | 8        | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 10       | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 12       | 1       |
| e0s8q070do39qcm87s3mktbfj6  | 8        | 1       | e0s8q070do39qcm87s3mktbfj6  | 10       | 1       | e0s8q070do39qcm87s3mktbfj6  | 12       | 1       |
| euc6ho5uqo8vputqo7u0jnihc1  | 8        | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 10       | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 12       | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 8        | 1       | frj8s3fd1po84uje7sftrh85m5  | 10       | 1       | frj8s3fd1po84uje7sftrh85m5  | 12       | 1       |
| gj0e3kmm5tuaasosvni dri1663 | 8        | 1       | gj0e3kmm5tuaasosvni dri1663 | 10       | 1       | gj0e3kmm5tuaasosvni dri1663 | 12       | 1       |
| i97epapdrhnh3sisupuju7orb0  | 8        | 1       | i97epapdrhnh3sisupuju7orb0  | 10       | 1       | i97epapdrhnh3sisupuju7orb0  | 12       | 1       |
| ij381i664l bmd2lrorlts4jn2  | 8        | 1       | ij381i664l bmd2lrorlts4jn2  | 10       | 1       | ij381i664l bmd2lrorlts4jn2  | 12       | 1       |
| msdnp47br71gdv11iuidgnrt43  | 8        | 0       | msdnp47br71gdv11iuidgnrt43  | 10       | 1       | msdnp47br71gdv11iuidgnrt43  | 12       | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 8        | 1       | nfnkmvflpcukq18pvdssds4r22  | 10       | 1       | nfnkmvflpcukq18pvdssds4r22  | 12       | 1       |
| o4ceqglln1nphd7c6kb4b93923  | 8        | 1       | o4ceqglln1nphd7c6kb4b93923  | 10       | 1       | o4ceqglln1nphd7c6kb4b93923  | 12       | 1       |
| oep8758jc2ae7983kml3qqupf2  | 8        | 1       | oep8758jc2ae7983kml3qqupf2  | 10       | 1       | oep8758jc2ae7983kml3qqupf2  | 12       | 1       |
| p3beelqe1gbd8tk2qic2hbfej5  | 8        | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 10       | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 12       | 1       |
| qt4orejugu7hqbd0m8j7r7j9d56 | 8        | 1       | qt4orejugu7hqbd0m8j7r7j9d56 | 10       | 1       | qt4orejugu7hqbd0m8j7r7j9d56 | 12       | 1       |
| rde6dsetg9tn721f8ef11giug5  | 8        | 1       | rde6dsetg9tn721f8ef11giug5  | 10       | 1       | rde6dsetg9tn721f8ef11giug5  | 12       | 1       |
| risfv5f45skj2k0hsjhb0cnp2   | 8        | 1       | risfv5f45skj2k0hsjhb0cnp2   | 10       | 1       | risfv5f45skj2k0hsjhb0cnp2   | 12       | 1       |
| umbkj3outn2q0u4ogi1nhm0d86  | 8        | 1       | umbkj3outn2q0u4ogi1nhm0d86  | 10       | 1       | umbkj3outn2q0u4ogi1nhm0d86  | 12       | 1       |
| vulccmpfd9kigtmm76qvekukl1  | 8        | 1       | vulccmpfd9kigtmm76qvekukl1  | 10       | 1       | vulccmpfd9kigtmm76qvekukl1  | 12       | 1       |

| Session                     | Question | Correct | Session                     | Question | Correct | Session                     | Question | Correct |
|-----------------------------|----------|---------|-----------------------------|----------|---------|-----------------------------|----------|---------|
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.3      | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.5      | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.7      | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 1.3      | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 1.5      | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 1.7      | 1       |
| 2mbikr6emk41b87kr46d68fh44  | 1.3      | 1       | 2mbikr6emk41b87kr46d68fh44  | 1.5      | 1       | 2mbikr6emk41b87kr46d68fh44  | 1.7      | 1       |
| 4pv1c0oi66fe06q2ck66l3o73   | 1.3      | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 1.5      | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 1.7      | 1       |
| 53tkvqq01r0jv30au4ikr2eb20  | 1.3      | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 1.5      | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 1.7      | 1       |
| 5f6po4ekr3bm57pik3md4tnub1  | 1.3      | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 1.5      | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 1.7      | 1       |
| 5n8u491vgs7r2pl76msom9s8v7  | 1.3      | 1       | 5n8u491vgs7r2pl76msom9s8v7  | 1.5      | 1       | 5n8u491vgs7r2pl76msom9s8v7  | 1.7      | 1       |
| 5rov1713rmqdn7mo4h715ln5n3  | 1.3      | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 1.5      | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 1.7      | 1       |
| 6rubc1hv5phadam9njf847an21  | 1.3      | 1       | 6rubc1hv5phadam9njf847an21  | 1.5      | 1       | 6rubc1hv5phadam9njf847an21  | 1.7      | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 1.3      | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 1.5      | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 1.7      | 1       |
| 7rj4aa2crif7lsrhson62ron41  | 1.3      | 1       | 7rj4aa2crif7lsrhson62ron41  | 1.5      | 1       | 7rj4aa2crif7lsrhson62ron41  | 1.7      | 1       |
| 8ejfs86cb24gludak26v5itop4  | 1.3      | 1       | 8ejfs86cb24gludak26v5itop4  | 1.5      | 1       | 8ejfs86cb24gludak26v5itop4  | 1.7      | 1       |
| 8ncopsrtca2mqh9qmulkc79t4   | 1.3      | 1       | 8ncopsrtca2mqh9qmulkc79t4   | 1.5      | 1       | 8ncopsrtca2mqh9qmulkc79t4   | 1.7      | 1       |
| 8o1in7idjghsf9va07uo6g1jk3  | 1.3      | 1       | 8o1in7idjghsf9va07uo6g1jk3  | 1.5      | 1       | 8o1in7idjghsf9va07uo6g1jk3  | 1.7      | 0       |
| 997f013mqp3p2h2qnp581cb76   | 1.3      | 1       | 997f013mqp3p2h2qnp581cb76   | 1.5      | 1       | 997f013mqp3p2h2qnp581cb76   | 1.7      | 1       |
| 9d1o1p49nbnuilumoutrjtcq2   | 1.3      | 1       | 9d1o1p49nbnuilumoutrjtcq2   | 1.5      | 1       | 9d1o1p49nbnuilumoutrjtcq2   | 1.7      | 1       |
| 9m4evllj7tm1q9icls815o40m4  | 1.3      | 1       | 9m4evllj7tm1q9icls815o40m4  | 1.5      | 1       | 9m4evllj7tm1q9icls815o40m4  | 1.7      | 1       |
| biabpokjnk132lu2keb9nhmt5   | 1.3      | 1       | biabpokjnk132lu2keb9nhmt5   | 1.5      | 1       | biabpokjnk132lu2keb9nhmt5   | 1.7      | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 1.3      | 1       | c7bmbkeriqjfr85sm3vcbth933  | 1.5      | 1       | c7bmbkeriqjfr85sm3vcbth933  | 1.7      | 1       |
| csli5u3ms3rtg8qht0rfihd6q3  | 1.3      | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 1.5      | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 1.7      | 1       |
| e0s8q070do39qcm87s3mktbfj6  | 1.3      | 1       | e0s8q070do39qcm87s3mktbfj6  | 1.5      | 1       | e0s8q070do39qcm87s3mktbfj6  | 1.7      | 1       |
| euc6ho5uqo8vputqo7u0jnihc1  | 1.3      | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 1.5      | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 1.7      | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 1.3      | 1       | frj8s3fd1po84uje7sftrh85m5  | 1.5      | 1       | frj8s3fd1po84uje7sftrh85m5  | 1.7      | 1       |
| gj0e3kmm5tuaasosvni dri1663 | 1.3      | 1       | gj0e3kmm5tuaasosvni dri1663 | 1.5      | 1       | gj0e3kmm5tuaasosvni dri1663 | 1.7      | 1       |
| i97epapdrhnh3sisupuju7orb0  | 1.3      | 1       | i97epapdrhnh3sisupuju7orb0  | 1.5      | 1       | i97epapdrhnh3sisupuju7orb0  | 1.7      | 1       |
| ij38i1664l bmd2lrorlts4jn2  | 1.3      | 1       | ij38i1664l bmd2lrorlts4jn2  | 1.5      | 1       | ij38i1664l bmd2lrorlts4jn2  | 1.7      | 1       |
| msdnp47br71gdv11iuidgnrt43  | 1.3      | 1       | msdnp47br71gdv11iuidgnrt43  | 1.5      | 1       | msdnp47br71gdv11iuidgnrt43  | 1.7      | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 1.3      | 1       | nfnkmvflpcukq18pvdssds4r22  | 1.5      | 1       | nfnkmvflpcukq18pvdssds4r22  | 1.7      | 1       |
| o4ceqglln1nphd7c6kb4b93923  | 1.3      | 1       | o4ceqglln1nphd7c6kb4b93923  | 1.5      | 1       | o4ceqglln1nphd7c6kb4b93923  | 1.7      | 1       |
| oep8758jc2ae7983kml3qqupf2  | 1.3      | 1       | oep8758jc2ae7983kml3qqupf2  | 1.5      | 1       | oep8758jc2ae7983kml3qqupf2  | 1.7      | 1       |
| p3beelqe1gbd8tk2qic2hbfej5  | 1.3      | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 1.5      | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 1.7      | 1       |
| qt4orejugu7hqbd0m8j7r79d56  | 1.3      | 1       | qt4orejugu7hqbd0m8j7r79d56  | 1.5      | 1       | qt4orejugu7hqbd0m8j7r79d56  | 1.7      | 1       |
| rde6dsetg9tn721f8ef11giug5  | 1.3      | 1       | rde6dsetg9tn721f8ef11giug5  | 1.5      | 1       | rde6dsetg9tn721f8ef11giug5  | 1.7      | 1       |
| risfv5f45sjk2k0hsjhb0cnp2   | 1.3      | 1       | risfv5f45sjk2k0hsjhb0cnp2   | 1.5      | 1       | risfv5f45sjk2k0hsjhb0cnp2   | 1.7      | 1       |
| umbkj3outn2q0u4og1inh0d86   | 1.3      | 1       | umbkj3outn2q0u4og1inh0d86   | 1.5      | 1       | umbkj3outn2q0u4og1inh0d86   | 1.7      | 1       |
| vulccmpfd9kigtmm76qvekukl1  | 1.3      | 1       | vulccmpfd9kigtmm76qvekukl1  | 1.5      | 1       | vulccmpfd9kigtmm76qvekukl1  | 1.7      | 1       |
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.4      | 1       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.6      | 0       | 0en70q9b8o6bn6bgfkhcf2i8h7  | 1.8      | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 1.4      | 0       | 1u07p8omppjb2oqrtkadvdn2v6  | 1.6      | 1       | 1u07p8omppjb2oqrtkadvdn2v6  | 1.8      | 1       |
| 2mbikr6emk41b87kr46d68fh44  | 1.4      | 1       | 2mbikr6emk41b87kr46d68fh44  | 1.6      | 1       | 2mbikr6emk41b87kr46d68fh44  | 1.8      | 1       |
| 4pv1c0oi66fe06q2ck66l3o73   | 1.4      | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 1.6      | 1       | 4pv1c0oi66fe06q2ck66l3o73   | 1.8      | 1       |
| 53tkvqq01r0jv30au4ikr2eb20  | 1.4      | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 1.6      | 1       | 53tkvqq01r0jv30au4ikr2eb20  | 1.8      | 1       |
| 5f6po4ekr3bm57pik3md4tnub1  | 1.4      | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 1.6      | 1       | 5f6po4ekr3bm57pik3md4tnub1  | 1.8      | 1       |
| 5n8u491vgs7r2pl76msom9s8v7  | 1.4      | 1       | 5n8u491vgs7r2pl76msom9s8v7  | 1.6      | 1       | 5n8u491vgs7r2pl76msom9s8v7  | 1.8      | 1       |
| 5rov1713rmqdn7mo4h715ln5n3  | 1.4      | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 1.6      | 1       | 5rov1713rmqdn7mo4h715ln5n3  | 1.8      | 1       |
| 6rubc1hv5phadam9njf847an21  | 1.4      | 1       | 6rubc1hv5phadam9njf847an21  | 1.6      | 1       | 6rubc1hv5phadam9njf847an21  | 1.8      | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 1.4      | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 1.6      | 1       | 74bi9bd7o2ukkvfso9m18jou11  | 1.8      | 1       |
| 7rj4aa2crif7lsrhson62ron41  | 1.4      | 1       | 7rj4aa2crif7lsrhson62ron41  | 1.6      | 1       | 7rj4aa2crif7lsrhson62ron41  | 1.8      | 1       |
| 8ejfs86cb24gludak26v5itop4  | 1.4      | 1       | 8ejfs86cb24gludak26v5itop4  | 1.6      | 0       | 8ejfs86cb24gludak26v5itop4  | 1.8      | 1       |
| 8ncopsrtca2mqh9qmulkc79t4   | 1.4      | 1       | 8ncopsrtca2mqh9qmulkc79t4   | 1.6      | 1       | 8ncopsrtca2mqh9qmulkc79t4   | 1.8      | 1       |
| 8o1in7idjghsf9va07uo6g1jk3  | 1.4      | 1       | 8o1in7idjghsf9va07uo6g1jk3  | 1.6      | 1       | 8o1in7idjghsf9va07uo6g1jk3  | 1.8      | 1       |
| 997f013mqp3p2h2qnp581cb76   | 1.4      | 1       | 997f013mqp3p2h2qnp581cb76   | 1.6      | 1       | 997f013mqp3p2h2qnp581cb76   | 1.8      | 1       |
| 9d1o1p49nbnuilumoutrjtcq2   | 1.4      | 1       | 9d1o1p49nbnuilumoutrjtcq2   | 1.6      | 1       | 9d1o1p49nbnuilumoutrjtcq2   | 1.8      | 1       |
| 9m4evllj7tm1q9icls815o40m4  | 1.4      | 1       | 9m4evllj7tm1q9icls815o40m4  | 1.6      | 0       | 9m4evllj7tm1q9icls815o40m4  | 1.8      | 1       |
| biabpokjnk132lu2keb9nhmt5   | 1.4      | 1       | biabpokjnk132lu2keb9nhmt5   | 1.6      | 1       | biabpokjnk132lu2keb9nhmt5   | 1.8      | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 1.4      | 0       | c7bmbkeriqjfr85sm3vcbth933  | 1.6      | 1       | c7bmbkeriqjfr85sm3vcbth933  | 1.8      | 1       |
| csli5u3ms3rtg8qht0rfihd6q3  | 1.4      | 1       | csli5u3ms3rtg8qht0rfihd6q3  | 1.6      | 0       | csli5u3ms3rtg8qht0rfihd6q3  | 1.8      | 1       |
| e0s8q070do39qcm87s3mktbfj6  | 1.4      | 1       | e0s8q070do39qcm87s3mktbfj6  | 1.6      | 1       | e0s8q070do39qcm87s3mktbfj6  | 1.8      | 1       |
| euc6ho5uqo8vputqo7u0jnihc1  | 1.4      | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 1.6      | 1       | euc6ho5uqo8vputqo7u0jnihc1  | 1.8      | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 1.4      | 1       | frj8s3fd1po84uje7sftrh85m5  | 1.6      | 1       | frj8s3fd1po84uje7sftrh85m5  | 1.8      | 1       |
| gj0e3kmm5tuaasosvni dri1663 | 1.4      | 1       | gj0e3kmm5tuaasosvni dri1663 | 1.6      | 1       | gj0e3kmm5tuaasosvni dri1663 | 1.8      | 1       |
| i97epapdrhnh3sisupuju7orb0  | 1.4      | 0       | i97epapdrhnh3sisupuju7orb0  | 1.6      | 1       | i97epapdrhnh3sisupuju7orb0  | 1.8      | 1       |
| ij38i1664l bmd2lrorlts4jn2  | 1.4      | 1       | ij38i1664l bmd2lrorlts4jn2  | 1.6      | 1       | ij38i1664l bmd2lrorlts4jn2  | 1.8      | 1       |
| msdnp47br71gdv11iuidgnrt43  | 1.4      | 0       | msdnp47br71gdv11iuidgnrt43  | 1.6      | 1       | msdnp47br71gdv11iuidgnrt43  | 1.8      | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 1.4      | 1       | nfnkmvflpcukq18pvdssds4r22  | 1.6      | 1       | nfnkmvflpcukq18pvdssds4r22  | 1.8      | 1       |
| o4ceqglln1nphd7c6kb4b93923  | 1.4      | 0       | o4ceqglln1nphd7c6kb4b93923  | 1.6      | 1       | o4ceqglln1nphd7c6kb4b93923  | 1.8      | 1       |
| oep8758jc2ae7983kml3qqupf2  | 1.4      | 1       | oep8758jc2ae7983kml3qqupf2  | 1.6      | 1       | oep8758jc2ae7983kml3qqupf2  | 1.8      | 1       |
| p3beelqe1gbd8tk2qic2hbfej5  | 1.4      | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 1.6      | 1       | p3beelqe1gbd8tk2qic2hbfej5  | 1.8      | 1       |
| qt4orejugu7hqbd0m8j7r79d56  | 1.4      | 1       | qt4orejugu7hqbd0m8j7r79d56  | 1.6      | 1       | qt4orejugu7hqbd0m8j7r79d56  | 1.8      | 1       |
| rde6dsetg9tn721f8ef11giug5  | 1.4      | 1       | rde6dsetg9tn721f8ef11giug5  | 1.6      | 1       | rde6dsetg9tn721f8ef11giug5  | 1.8      | 1       |
| risfv5f45sjk2k0hsjhb0cnp2   | 1.4      | 1       | risfv5f45sjk2k0hsjhb0cnp2   | 1.6      | 1       | risfv5f45sjk2k0hsjhb0cnp2   | 1.8      | 1       |
| umbkj3outn2q0u4og1inh0d86   | 1.4      | 1       | umbkj3outn2q0u4og1inh0d86   | 1.6      | 1       | umbkj3outn2q0u4og1inh0d86   | 1.8      | 1       |
| vulccmpfd9kigtmm76qvekukl1  | 1.4      | 1       | vulccmpfd9kigtmm76qvekukl1  | 1.6      | 1       | vulccmpfd9kigtmm76qvekukl1  | 1.8      | 1       |

| Session                     | Question | Correct |
|-----------------------------|----------|---------|
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 19       | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 19       | 1       |
| 2mbikr6emk41b87kr46d68fh44  | 19       | 1       |
| 4pv1c0oi66fe06q2ck66lg3o73  | 19       | 1       |
| 53tkvqq01r0jv30au4ikr2eb20  | 19       | 1       |
| 5f6po4ekr3bm57pik3md4tnub1  | 19       | 1       |
| 5n8u49lvgs7r2pl76msom9s8v7  | 19       | 1       |
| 5rov1713rmqdn7mo4h7l5ln5n3  | 19       | 1       |
| 6rubc1hv5phadum9njf847an21  | 19       | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 19       | 1       |
| 7rj4aa2crif7lshrson62ron41  | 19       | 1       |
| 8ejfs86cb24gludak26v5itop4  | 19       | 1       |
| 8ncopsrtca2mqqh9qmulkc79t4  | 19       | 1       |
| 8o1in7idjhsf9va07uo6g1jk3   | 19       | 1       |
| 997f013mqp3p2h2qppm58lcb76  | 19       | 1       |
| 9d1o1p49bnuilumoutrjtcq2    | 19       | 1       |
| 9m4evllj7tm1q9icls8l5o40m4  | 19       | 1       |
| biabpokjnk132lu2keb9nhmt5   | 19       | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 19       | 1       |
| cs1i5u3ms3rtg8qhtOrfihd6q3  | 19       | 1       |
| e0s8q070do39qcm87s3mktbfj6  | 19       | 1       |
| euc6ho5uqo8vputqo7u0jnihc1  | 19       | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 19       | 1       |
| gj0e3kmm5tuaasosvndri1663   | 19       | 1       |
| i97epapdrhnh3sisupuju7orb0  | 19       | 1       |
| ij38i664l bmd2lrorlts4jnj2  | 19       | 1       |
| msdnp47br7lgdv11iu1dgnrt43  | 19       | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 19       | 1       |
| o4ceqglln1nphd7c6kb4b93923  | 19       | 1       |
| oep8758jc2ae7983kml3qqupf2  | 19       | 1       |
| p3beelqe1gbd8tk2qic2hbfje5  | 19       | 1       |
| qt4orejugu7hqbd0m8jxr7j9d56 | 19       | 1       |
| rde6dsetg9tn721f8ef11giug5  | 19       | 1       |
| risfv5f45sjk2k0hsjhb0cnp2   | 19       | 1       |
| umbkj3outn2q0u4ogi1nhm0d86  | 19       | 1       |
| vulccmpfd9kigtmm76qvekukl1  | 19       | 1       |
| 0en70q9b8o6bn6bgfkhcf2i8h7  | 20       | 1       |
| 1u07p8omppjb2oqrtkadvdn2v6  | 20       | 0       |
| 2mbikr6emk41b87kr46d68fh44  | 20       | 1       |
| 4pv1c0oi66fe06q2ck66lg3o73  | 20       | 1       |
| 53tkvqq01r0jv30au4ikr2eb20  | 20       | 0       |
| 5f6po4ekr3bm57pik3md4tnub1  | 20       | 1       |
| 5n8u49lvgs7r2pl76msom9s8v7  | 20       | 1       |
| 5rov1713rmqdn7mo4h7l5ln5n3  | 20       | 0       |
| 6rubc1hv5phadum9njf847an21  | 20       | 1       |
| 74bi9bd7o2ukkvfso9m18jou11  | 20       | 1       |
| 7rj4aa2crif7lshrson62ron41  | 20       | 1       |
| 8ejfs86cb24gludak26v5itop4  | 20       | 1       |
| 8ncopsrtca2mqqh9qmulkc79t4  | 20       | 1       |
| 8o1in7idjhsf9va07uo6g1jk3   | 20       | 0       |
| 997f013mqp3p2h2qppm58lcb76  | 20       | 1       |
| 9d1o1p49bnuilumoutrjtcq2    | 20       | 1       |
| 9m4evllj7tm1q9icls8l5o40m4  | 20       | 1       |
| biabpokjnk132lu2keb9nhmt5   | 20       | 1       |
| c7bmbkeriqjfr85sm3vcbth933  | 20       | 0       |
| cs1i5u3ms3rtg8qhtOrfihd6q3  | 20       | 0       |
| e0s8q070do39qcm87s3mktbfj6  | 20       | 1       |
| euc6ho5uqo8vputqo7u0jnihc1  | 20       | 1       |
| frj8s3fd1po84uje7sftrh85m5  | 20       | 1       |
| gj0e3kmm5tuaasosvndri1663   | 20       | 1       |
| i97epapdrhnh3sisupuju7orb0  | 20       | 1       |
| ij38i664l bmd2lrorlts4jnj2  | 20       | 1       |
| msdnp47br7lgdv11iu1dgnrt43  | 20       | 1       |
| nfnkmvflpcukq18pvdssds4r22  | 20       | 1       |
| o4ceqglln1nphd7c6kb4b93923  | 20       | 1       |
| oep8758jc2ae7983kml3qqupf2  | 20       | 1       |
| p3beelqe1gbd8tk2qic2hbfje5  | 20       | 0       |
| qt4orejugu7hqbd0m8jxr7j9d56 | 20       | 1       |
| rde6dsetg9tn721f8ef11giug5  | 20       | 1       |
| risfv5f45sjk2k0hsjhb0cnp2   | 20       | 1       |
| umbkj3outn2q0u4ogi1nhm0d86  | 20       | 1       |
| vulccmpfd9kigtmm76qvekukl1  | 20       | 1       |