



Brock University

Department of Computer Science

Procedural 3D Texture Synthesis Using Genetic Programming

Adam Hewgill and Brian J. Ross
Technical Report # CS-03-06
April 2003

Brock University
Department of Computer Science
St. Catharines, Ontario
Canada L2S 3A1
www.cosc.brocku.ca

Procedural 3D Texture Synthesis Using Genetic Programming

Adam Hewgill Brian J. Ross¹

*Brock University, Dept. of Computer Science
St. Catharines, Ontario, Canada L2S 3A1*

Abstract

The automatic synthesis of procedural textures for 3D surfaces using genetic programming is investigated. Genetic algorithms employ a search strategy inspired by Darwinian natural evolution. Genetic programming uses genetic algorithms on tree structures, which are interpretable as computer programs or mathematical formulae. We use a texture generation language as a target language for genetic programming, and then use it to evolve textures having particular characteristics of interest. The texture generation language used here includes operators useful for texture creation, for example, mathematical operators, and colour and noise functions. In order to be practical for 3D model rendering, the language includes primitives that access surface information for the point being rendered, such as coordinates values, normal vectors, and surface gradients. A variety of experiments successfully generated procedural textures that displayed visual characteristics similar to the target textures used during training.

Key words: procedural textures, evolution, genetic programming

Email addresses: ah00bt@cosc.brocku.ca (Adam Hewgill),
bross@cosc.brocku.ca (Brian J. Ross).

URL: <http://www.cosc.brocku.ca/~bross/> (Brian J. Ross).

¹ Corresponding author. Equal authorship implied.

1 INTRODUCTION

Procedural textures are an integral part of contemporary rendering technology [1] [12]. Procedural textures can convincingly model a variety of natural phenomena, such as wood, stone, and terrain, as well as innumerable unnatural effects. Moreover, procedural textures apply rendered results seamlessly over surfaces, without tiling or other repetitive artifacts. This makes them ideal for photorealistic rendering.

One disadvantage of procedural textures is their technical complexity. Mathematical and procedural models of well-known effects, such as stone, wood, and terrain, have been carefully engineered and hand-crafted. Typically, parameterized versions of these pre-defined effects are included in rendering tools, along with suitable interfaces for user-manipulation of the parameters. The derivation of a new procedural texture can be a challenging task. If the user has a visual effect in mind, the derivation of a corresponding procedural texture that produces this effect will depend upon the user's proficiency and experience in mathematical texture modeling. Even for experts, trial-and-error will be the norm. Hence, the invention of computer-based tools which can automatically derive procedural textures is a worthy goal.

This paper investigates the automatic derivation of procedural textures for 3D surfaces using genetic programming technology. Genetic programming is a machine learning paradigm inspired by Darwinian evolution. Populations of computer programs are bred until a program solving some problem of interest is eventually evolved. We use genetic programming to evolve populations of 3D texture formulae. The approach is similar to previous work in 2D procedural texture evolution [3][9][13]. The main enhancement necessary for evolving effective 3D textures is to incorporate surface information into texture formulae. The texture language implemented in our experiments includes provision for surface coordinates, surface orientation (normal direction), and surface gradient (normal deviation). As in 2D texture evolution, a target texture is used for training purposes. Surface points on an example model rendered with the target texture are manually sampled. These sample points are then used as a training set by the genetic programming system. The goal is to evolve a 3D texture whose colour rendering on 3D surface components is as similar as possible to that in the target texture.

Section 2 reviews relevant work in texture evolution. An overview of genetic algorithms and genetic programming is given in Section 3. System and experimental details are discussed in Section 4. Some results are given in Section 5. Section 6 gives some concluding discussion.

2 Texture Evolution

Genetic algorithms have been used for interactively exploring procedural textures [8] [10] [11]. A number of commercial products have used this approach; Kai's Power Tools and Alien Skin Textureshop are two examples. These systems rely on the user to interactively participate in the evolution process. The scheme used is as follows. A population of candidate textures is displayed to the user. He or she selects the texture(s) of interest, to be mutated or recombined into the next set of textures. The user also controls a mutation rate parameter, which determines the degree to which new textures differ from current textures. The mutation rate is usually large at the beginning of a session, and then is gradually reduced when a texture of interest is discovered. In essence, the user becomes an oracle or "fitness evaluation function" for the duration of the search. Although this procedure is a useful interactive tool for texture exploration, a disadvantage is that it is not automated – the user plays the pivotal role in guiding evolution at all stages. The user will suffer fatigue if he or she must evaluate many hundreds of textures.

Fully automated texture generation using genetic algorithms has been investigated [3][9][13]. In order to automate the search, genetic algorithms must select candidate solutions based upon their fitness scores, which are indications of how well the textures conform to some set of desired visual characteristics. The challenge in automated texture evolution is to derive an effective fitness function to rank candidate textures based on their visual characteristics. Currently, it is unknown how to derive a formula that rates aesthetic qualities of a candidate texture. Instead, a suite of feature tests are used, which rate rudimentary visual characteristics of a texture, such as colour distribution, luminosity, and shape. Although each feature test by itself is a crude measurement of an image's characteristics, they often produce useful evaluations of textures when combined together. During a session, the feature test scores are compared to the corresponding feature scores of a target texture image. The goal for evolution is to derive a procedural texture whose visual characteristics closely match those of the target texture image, as indicated by the proximity of feature test scores between the target and candidate textures. Note that it is not realistic or desired for the genetic algorithm to derive a texture identical to the target texture, but rather, to evolve one that is similar in flavour to the target.

Ibrahim's Genshade system evolves Renderman shaders [3]. High-level Renderman shader expressions are denoted as dataflow graphs. Feature tests such as luminosity (brightness), chromaticity (colour), and wavelet analysis (shape) are performed. Evolution can proceed interactive with user guidance, or automatically with feature test evaluation. Multiple target textures can be used as well. The results show effective evolution of Renderman shaders that match

target textures generated by Renderman itself.

The Gentropy system by Wiens and Ross uses a similar approach as Genshade [13]. Here, textures are formulated as tree expressions, and the texture primitives are basic arithmetic, colour, and noise generators. A suite of feature tests are used, such as luminosity, colour, and wavelet tests. Flexible user-defined topologies of multiple populations and testing strategies are possible. Multiple target textures can also be used, in which different textures can be assigned to influence shape or colour. Gentropy is strictly for automated texture evolution, and no user interaction is possible at present.

Both Genshade and Gentropy treat the texture search space as a single objective problem, in which various feature test scores are combined arithmetically to define a single feature test score. With such a search space, the genetic algorithm's task is to find the optimal point – the highest or best score. This single-objective characterization of an inherently multi-objective problem such as texture analysis requires an arbitrary reconciliation of how differing feature tests relate to the overall score. For example, one simple scheme is to merely give all feature tests an equal weight in the overall final score. This is unsatisfactory in many instances, because one feature score may prematurely dominate the others, which results in a biased solution.

The multiple feature tests used by Genshade and Gentropy are more naturally denoted as a multi-objective problem (MOP). Work by Ross and Zhu consider texture evolution to be a multi-objective search problem [9]. A MOP is characterized by having two or more features, which may not be related in a straight-forward, linear fashion. In the Ross and Zhu system, the Gentropy engine is enhanced so that each feature test is treated as a separate dimension of a multi-dimensional search problem, and texture scores are rated using a Pareto ranking technique [2]. Pareto ranking means that colour and shape scores are never combined, but remain independent during the run. An advantage of this is that complex multi-population topologies with clever, customized feature testing strategies are unnecessary. A practical outcome of the MOP approach is that results competitive with the Gentropy system were obtained much more easily and efficiently (often 10 times faster).

The research in [3][9][13] is primarily focussed on the evolution of 2D procedural textures. Although these textures can be rendered onto 3D surfaces, the results are usually unsatisfactory, since such textures do not consider morphological features of the surface. This results in surface textures which are uniform across the entire 3D surface, and which do not react to surface characteristics such as location or orientation in 3-space.

3 Evolutionary Computation

3.1 Genetic algorithms

Evolutionary algorithms are search techniques inspired by Darwinian’s theory of evolution. One such evolutionary algorithm is the genetic algorithm (GA), which has proven very successful in finding good solutions for difficult real-world problems [5] (see Figure 1). In a GA, a population of candidate solutions is maintained. Each individual in the population is represented by a chromosome, which encodes some candidate solution to the problem at hand. Additionally, each individual is assigned a fitness score, which measures the quality of the solution encoded by the individual. These fitness scores are used to probabilistically select individuals for reproduction, in the sense that fitter individuals will be more likely to be selected than those less fit. This parallels Darwinian evolution’s “survival of the fittest”, in that fit individuals survive and procreate, while weak individuals die off and become extinct.

A defining characteristic of a genetic algorithm is the means by which new candidate solutions are generated. The most important reproduction operation in a GA is crossover. Crossover is inspired by natural sexual reproduction, and is the means by which parental traits are inherited by offspring. Basically, to perform crossover, two parents are selected based on their fitness, and their chromosomes are randomly split and merged together to form two offspring. For example, consider two individuals with chromosomes $ABCDEFGG$ and $tuvwxyz$ respectively. A single-point crossover operation finds some random split point in the 7-gene chromosome. If this point is 3, then the offspring become:

$$\begin{array}{ccc} \underline{ABC}DEFG & \implies & \underline{tw}DEFG \\ \underline{tuvw}xyz & & \underline{ABC}wxyz \end{array}$$

The other important genetic reproduction operator is mutation. This involves randomly changing some random gene in a chromosome. For example, the individual $ABCDEFGG$ may become $ABxDEFG$ by a random mutation of the gene C to x . Typically, mutation plays a lesser role than crossover, and is used more sparingly.

It is worth reiterating the importance of fitness-based selection in the above discussion. In order for a genetic algorithm to be successful, individuals used for reproduction must be selected based upon their relative fitness or strength in solving some problem of interest. There are a number of schemes for performing fitness-based selection. The one used in this paper is the *tournament selection*. Here, a set of 2 or more individuals are randomly selected from the population. Then the individual with the highest fitness score is retained as

a parent for reproduction. This tournament selection is performed once for mutation, and twice for crossover (once for each parent).

3.2 Genetic programming

Genetic programming (GP) is a specialized class of genetic algorithm, in which individuals in the population are interpretable as computer programs [4]. One of genetic programming's strengths as a paradigm of evolutionary computation is its applicability to a wide variety of problems. Any application that is best solved via an algorithm or mathematical formula is a good candidate for genetic programming. This differs from a vanilla genetic algorithm that uses bitstrings as chromosomes, in which such strings might require complex and clever coding schemes in order to be applied to problems with dynamically changing requirements.

This paper uses the common tree-based representation of programs. In this representation, internal nodes denote function calls, and nodes below them denote arguments (Figure 2). In order for this representation to be usable, it must be ensured that all trees processed are both syntactically correct, and executable without errors. Executability or *closure* is maintained by ensuring that every function used in the program will execute without errors on any supplied input data. For example, a closed division operator will not invoke an error condition if a division by zero is attempted, but instead will return some predefined value as an answer.

Syntactic correctness of programs is required if programs are to execute in a correct and predictable fashion. Firstly, random trees as created in the initial population must be syntactically sound. To do this, the random tree generator must supply any function in a tree with its required arguments. Reproduction must also maintain syntactic correctness. After two parents are selected for crossover, random nodes are selected in each parent's tree. The subtrees defined at these nodes are then swapped to create the offspring. This maintains syntactic correctness. Mutation involves replacing a randomly selected subtree with a randomly generated subtree. A result of these tree-based reproduction operators is that chromosomes are variable-sized, and can grow to unbounded sizes if not checked. Tree-depth limits are typically used in GP runs.

4 Experiment

4.1 Graphics environment

A fundamental characteristic of procedural textures is that they define a colour for all points in the coordinate space. This makes them ideal for ray-traced environments, in which textures can be calculated for all visible hit points on model surfaces. On the other hand, they are less useful for rasterized graphics, in which lighting is computed for vertices only, and polygon surface shading is interpolated via Gouraud or Phong shading, mixed perhaps with texture maps. As a consequence, a ray-tracing environment is used in this research. The intention is for the evolved procedural texture formulae to be supplied to a ray-tracer for rendering purposes. However, actual rendering is not performed during texture evolution, but rather, texture formulae will be evaluated for defined sets of example points. Other well-known ray-tracer effects, such as local lighting, reflection, transmission, and shadows, are not considered nor implemented.

Theoretically, any surface texture is definable by some suitable procedural texture. Naturally, the complexity of the corresponding texture formula or algorithm can vary considerably, depending on the match between texture primitives and the target texture desired. At a minimum, a 3-D texture requires the XYZ coordinates of a surface point in order to calculate that point's RGB value:

$$f : (x, y, z) \rightarrow (r, g, b)$$

Such textures are defined entirely by the position of surface points in the coordinate space. For some textures, such as stone and marble, this is satisfactory and adequate. More complex effects, such as the natural rendering of a mountain surface, with a snow-capped peak, rocky cliff, and green base, can also be handled to an extent by these formulae. To do so, the height of the model is influential in producing the appropriate effect as seen in the model's altitude. Often, however, we require textures that conform to additional local shape characteristics of a model. For example, a mountain surface might turn rocky wherever there is a steep cliff, which might occur at various altitudes. The steepness of the cliff might be modeled as a function of the surface normal. In other instances, we might like a texture which changes when there are abrupt discontinuities or creases on a surface. Such features are readily denoted via gradients (the degree to which the surface normal changes).

We used the following model surface information in various experiments:

- X, Y, Z coordinate: This will be the coordinate of the hit point on the model surface as computed during ray-tracing.
- Surface normal: the normal of a hit point.
- Interpolated mesh normal: This is Phong-style normal interpolation on a polygonal mesh. Shared vertex normals are interpolated from surrounding polygons. Then a hit point normal is interpolated from the surrounding vertex normals.
- Surface gradient: This is a scalar value representing the degree to which normals on a local area of a surface are changing. First, the interpolated mesh normal N^i for a hit point is calculated (see above). Then the difference in unsigned magnitude between it and the real surface normal N of the plane on which it resides is found. These differences are then averaged:

$$(|N_x^i - N_x| + |N_y^i - N_y| + |N_z^i - N_z|)/3$$

The above surface parameter definitions are used for generating training data for target textures. An interactive texture sampling editor is used for this purpose. A 3-D model with an applied bitmapped texture is read into the system. The user then manually selects surface points to be used as training data. For all points selected, a record of the above surface information is generated and saved in a table, along with the corresponding RGB colour used to render that sample. This table is saved in a text file, for use by the genetic programming system.

The inclusion of normals and gradients in textures results in considerable differences in texture characteristics. Without normals and gradients, surface detail disappears for all but the simplest models (Figure 3, a). Adding normals into texture formulae permits surface detail to arise without the use of lighting (Figure 3, b), and gradients introduce an additional dimension of surface detail (Figure 3, c). Gradient-based textures create shiny, (false) specular highlights, and are often metallic and iridescent in appearance. As is shown in Section 5.4, normals and gradients can also be useful during training.

4.2 Evaluation strategy

Once a training set of texture information is sampled by the user, we wish to evolve a procedural texture that exhibits the texture characteristics as encoded in the training set. Given a candidate texture formula, a requirement will be to evaluate how well its rendered texture corresponds to the sampled target texture. Such evaluations take the form of fitness scores, in which a high score means a high (strong) fitness.

A simple evaluation strategy will be used to determine the fitness score of a texture formula. The surface parameters of each sample point in the training file will be made available to the procedural texture. The texture formula is then interpreted on a training sample. It will make free use of the example surface parameters whenever they are required. Eventually, the formula computes an RGB colour for that point. The distance in RGB space between the actual rendered colour RGB_a and the sampled target colour RGB_s is determined. This distance is then tallied for all k samples in the test set:

$$Fitness = \sum_{i=1}^k dist(RGB_a, RGB_s)$$

where

$$dist(RGB_1, RGB_2) = \sqrt{(R_1 - R_2)^2 + (B_1 - B_2)^2 + (G_1 - G_2)^2}.$$

4.3 Texture language

The genetic programming system used here is based on Koza’s Lisp-based genetic programming model, in which programs take the form of symbolic expressions or *s-expressions* [4]. An s-expression is a list or atom in the Lisp language. An example expression in Lisp notation is:

$$(/ (+ 6 6) 4)$$

This denotes the arithmetic expression $(6 + 6)/4$. The mapping between s-expressions and trees is straight-forward: the first element in a list is the function (subtree root), its arguments are branches, and atoms are leaves.

The strongly-typed lilGP 1.1 system is used as the GP platform for the experiments [14]. This is a C-based GP system, which implements Koza’s tree-based GP paradigm. It is embellished with strong-typing, which means that expressions can be assigned to return designated data types [6].

Figure 4 shows a Backus-Naur Form grammar of the texture language. Note that expressions are denoted in standard algebraic notation, rather than as Lisp s-expressions. There are two data types in the language – rgb vectors (RGB) and floating point values (F). The RGB operators consist of either the current 3-D coordinate of the point being rendered, the current surface normal of the point being rendered, or an RGB vector *rgbvec* constructed from the values of three floating point expressions. Note that the RGB-space is modeled with clamped floating-point values between 0.0 and 1.0.

The rest of the grammar defines floating point expressions. Floating point terminals can be basic surface information for the point being rendered, such as one of its coordinate values (x, y, z) , normal directions (n_x, n_y, n_z) or surface gradient (*diff*). The other floating point terminal, *ephem*, is an ephemeral random constant. This is a constant value that is initialized with a random value between 0.0 and 1.0, but then retains its initialized value throughout its lifetime during the evolution process. In other words, it does *not* denote a call to a random number generator.

The floating point functions include sine, cosine, minimum, maximum, and conventional arithmetic operators. The divide operator was omitted, due to its tendency to create extremely large or small values, which in turn nullify the utility of many formula. The *avg* function finds the average between two expressions, while *lum* finds the luminosity (average channel value) of an RGB vector. The *if* expression interprets its first floating point argument. If that expression is greater than 0.5 in value, then the value of the second argument is computed and returned as a result. Otherwise the value of the third argument is returned. Finally, *noise* generates Perlin noise computed to the 4th harmonic [7].

Because the texture language is strongly typed, tree generation and reproduction will maintain legal data typing at all times. During generation, any tree generated must adhere to the typing conventions of each function shown in the grammar in Figure 4. For example, if a function requires an RGB argument, then an RGB function or terminal will be randomly selected. When applying crossover, the data type of the root of each subtree to be swapped must also match.

4.4 GP parameters

Table 1 lists the genetic programming parameters common to all the experiments. All the experiments used a population of size 1000 running for at least 200 generations. This was done for 10 runs per experiment, where each run has a new random number seed. Therefore, each run processes at least 200,000 textures, and each experiment at least 2 million textures. This processing was greatly aided by the use of a 16-CPU Silicon Graphics Origin 2000 server, which permits 10 runs to execute concurrently.

The initial randomized population is created using Koza's ramped half&half tree generation strategy from [4]. Half the trees are *grow trees*, in which each randomly generated node has an equal chance of being a function (internal node) or terminal (leaf), up to a maximum depth for the tree. The other trees are *full trees*, where nodes are leaf nodes only when the maximum depth of

the tree has been reached. Hence grow trees tend to be more asymmetric and smaller in size than the larger, bushier full trees. The idea of *ramping* means that maximum tree sizes are iterated between 5 to 10 levels deep for both grow and full trees. Ramped tree generation proceeds until the population is filled.

During evolution, crossover is used to create 90% of a new population, and mutation is used for the remainder. Trees can never exceed a maximum depth of 17 levels. If they do, the reproduction operator is tried again with new parents. Parent(s) are selected using a tournament selection with a tournament size of 5. This tournament size creates a fairly strong selection pressure.

5 Results

Results from selected experiments are now given. All the experiments share the genetic programming parameters of Section 4.4. Experiments differ with respect to the composition of training sets, the subset of the texture language used as a target language for evolved results, and in a few cases, the number of generations. In the remainder of the section, alternate models are rendered with solution textures by rescaling their coordinate extents to be those of the original training model.

5.1 Primary Colour Cube

This first example uses a simple training set (Figure 5). Six sample points are selected from centers of sides of a cube. These points are assigned the colours red, green, blue, grey, white, and yellow. As is done in all experiments, the fitness function evaluates how closely a texture formula renders these points with colours near their sampled target colours. This does not imply that the entire side of the cube will necessarily be that colour, but only the single center point sampled for training. This indeed can be seen in the cube images in the figure. The gradient and noise primitives are removed from the texture language for this run.

Figure 6 shows the evolved texture formula used in column 3 of Figure 5. The size and complexity of this formula is a fairly typical result from genetic programming. There is room for simplification in the formula. For example, the term $(\min 0.36341 0.13969)$ found 15 times in the formula can be replaced with the constant 0.13969 . Such extraneous expressions are called *intron code* or *program bloat*, and have no deleterious effect on the quality of a solution, other than making interpretation slower than necessary.

5.2 *Terrain Texture*

Figure 7 shows the results of evolved textures suitable for a mountainous landscape. The training model (top row) consists of 97 points sampled from a polygonal mesh mountain. The intension is for the mountain to have a white snowcap, grey cliffs, and a green base.

5.3 *Clothing Texture*

The next example is a bit offbeat. A female figure is used as a surface for training (Figure 8). A total of 1327 sampled points are used. The idea is that the procedural texture will colour the woman by giving her blonde hair, a green tank-top shirt, purple pants, black shoes, and pink face and arms. This is a complex task for a mathematical texture formula, given that the training file does not have high-level information about model components (arms, legs,...), but only the coordinate, normal, and gradient information of each sample point. The results after 600 generations are shown in Figure 9. The terrain images use the last clothing texture (row 2, column 2).

Figure 10 shows the fitness progress of the run that evolved the texture in row 2, column 2 of Figure 9. As can be seen, the fitness of the best individual and population average shows steady progress through the entire 600 generations of the run. This shows that 600 generations is not excessive. In fact, additional generations would likely give further progress, based on the trend in this graph.

5.4 *Normal and Gradient Specialized Texture*

Figure 11 shows the evolution of a texture that captures a particular surface orientation and shape for a model. The training model (column 1) consists of 5 steep hills on a flat surface. The tip of each hill is to be coloured red. The tip area has a high gradient, and variable normal. The rest of the hill below the tip is green. The flat base from which the hills rise is white. A total of 115 training points were used.

5.5 *Miscellaneous Results*

Figure 12 shows an experiment in which a terrain-oriented training set similar to that in Section 5.2 is used. After completing the sampling of the training set, however, we swapped the blue channel and gradient values. It is difficult

to intuit the effect of this new interpretation of data. The runs yielded some unexpected and interesting textures, and models tended to appear made of translucent minerals.

The final experiment uses the following training set. A total of 182 sample points are used upon a cube (Figure 13). The cube is blue, with 3 coloured bands (yellow, purple, tan) wrapping around the cube in the middle of each face on each of the XY, XZ and YZ planes. The 8 vertices have shared normals, interpolated from the faces surrounding each vertex. Hence the normals do not indicate the planar normals of each face. Since this is a particularly challenging texture to evolve, we let the system run for a total of 600 generations, rather than the usual 200 generations. Some results are shown in Figure 14. The results of training (column 1) are unexceptional. This is likely due to descriptive shortcomings in the texture language, which is clearly lacking in primitives that are adaptable to this particular target texture. Nevertheless, the results of applying the solution texture to other models are interesting, especially given that the overall colours used in the training set were prevalent.

6 Conclusion

This paper reports a first investigation into the suitability of using evolutionary computation to automatically synthesize 3-D procedural textures. The results of these experiments are positive and promising. Our scoring method – positive example matching – is very simple. This makes the results even more impressive, since 3-D procedural texture spaces can be complex and high-dimensional. The randomness and chaos that is an inherent part of evolutionary computation is a distinct advantage in texture generation, as it lends an element of invention and surprise to the texture synthesis process. Admittedly, some of the target textures used here are simple enough that handwritten solutions could be derived. The important point, however, is that no manual derivation of procedural textures was necessary – genetic programming obtained results automatically from training data. Obtaining samples for training is a much simpler task than deriving texture formulae.

There are a number of directions for improving this research. Firstly, the texture language used in this paper is very rudimentary, having only a basic set of arithmetic and RGB operators, along with a basic noise primitive. More complex textures will arise with a language that has higher-level texture generating primitives. For example, the use of Renderman shaders as done in Genshade would immediately result in more complex texture results [3]. The language could also be supplemented with additional 3-D model information. For example, a notion of model hierarchy might be introduced into the language. This would permit the texture to incorporate which portion of the

model hierarchy a texture is to be rendered upon. This would yield textures that are tuned more intimately to the model's structure.

More work needs to address the issue of fitness evaluation. Although the positive example matching approach used here is simple to implement, and often yields acceptable results, it could use improvement. A fundamental problem with example scoring is that the results are strongly dependent on the fairness of the example set. For example, if the majority of examples are blue, then fitness evaluation will naturally be biased towards formulae yielding blue colours. To overcome this bias, the user needs to balance the example set in a way that different features are adequately populated within the set.

Another problem with example scoring is that positive example matching is unsuitable for evolving noisy textures. The nature of noise is that a target pixel satisfying various surface characteristics may have substantially variable colour depending on the chaotic nature of the noisy texture space itself. Unfortunately, the use of positive example matching is decidedly prejudiced against noise, because of the chaotic colouring that occurs when noise primitives are present in a texture formula. For example, a noisy area of texture space may generate colours within some distribution. Fitness scoring, however, requires a strict match with the colour designated within the example. This means that the existence of a noise function will detract from fitness, and so evolution will prefer non-noisy formulae with more deterministic rendering behaviour. Our experience is that noise primitives usually tend to disappear from the population during early generations.

An alternative fitness evaluation technique that would not be biased against noise would be to incorporate probabilistic matching of candidate and target examples. For example, one could ascribe to an example point a probabilistic colour distribution. This would presume that a number of examples would be included in this probabilistic colour set, so that the resulting texture's colour distribution would closely align itself to the example set.

We found that the most pleasing aesthetic results do not necessarily correlate with the strongest fitness scores. Therefore, although fitness guides evolution, it is clear that artistic applications such as this one require human evaluation to determine the appropriateness of overall solutions. Although automated 3-D texture evolution is a challenging and worthy topic, it is interesting to consider the implementation of semi-automated texture synthesis systems that incorporate a dynamic interplay of automated and user-directed evaluation. Such tools might in fact turn out to be the most practical.

Acknowledgement: This research is supported by NSERC Operating Grant 138467-1998.

References

- [1] D.S. Ebert, F.K. Musgrave, D. Peachey, K. Perlin, and S. Worley. *Texturing and Modeling: a Procedural Approach*. Academic Press, 2 edition, 1998.
- [2] D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, 1989.
- [3] A.E.M. Ibrahim. *GenShade: an Evolutionary Approach to Automatic and Interactive Procedural Texture Generation*. PhD thesis, Texas A&M University, December 1998.
- [4] J.R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, 1992.
- [5] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1996.
- [6] D.J. Montana. Strongly Typed Genetic Programming. *Evolutionary Computation*, 3(2):199–230, 1995.
- [7] K. Perlin. An Image Synthesizer. *Computer Graphics*, 19(3), 1985.
- [8] S. Rooke. Eons of Genetically Evolved Algorithmic Images. In P.J. Bentley and D.W. Corne, editors, *Creative Evolutionary Systems*, pages 330–365. Morgan Kaufmann, 2002.
- [9] B.J. Ross and H. Zhu. Procedural Texture Evolution Using Multiobjective Optimization. Technical Report CS-02-18, Brock University, Dept. of Computer Science, July 2002.
- [10] A. Rowbottom. Evolutionary Art and Form. In P.J. Bentley, editor, *Evolutionary Design by Computers*, pages 330–365. Morgan Kaufmann, 1999.
- [11] K. Sims. Interactive evolution of equations for procedural models. *The Visual Computer*, 9:466–476, 1993.
- [12] A. Watt and M. Watt. *Advanced Animation and Rendering Techniques: Theory and Practice*. ACM Press, 1992.
- [13] A.L. Wiens and B.J. Ross. Gentropy: Evolutionary 2D Texture Generation. *Computers and Graphics Journal*, 26(1):75–88, February 2002.
- [14] D. Zongker and B. Punch. *lil-gp 1.0 User's Manual*. Dept. of Computer Science, Michigan State University, 1995.

1. Initialize the population with random chromosomes.
Rate the fitness of the individuals in the initial population.
2. Repeat until solution found OR maximum generation reached:
 - i) Repeat until new population generated:
 - a) Select 2 individuals based on their fitness, and apply crossover to generate two offspring. Add them to new population.
 - b) Select an individual based on fitness, and apply mutation. Add mutant to new population.
 - ii) Rate the fitness of individuals in new population.

Fig. 1. Example genetic algorithm




Fig. 2. Example tree and corresponding expression




Fig. 3. Texture language effects. (a) Basic. (b) Normals. (c) Normals and gradients.

$$\begin{aligned}
RGB & ::= RGB_{term} \mid RGB_{func} \\
RGB_{term} & ::= v_{xyz} \mid n_{xyz} \\
RGB_{func} & ::= rgbvec(F, F, F) \\
F & ::= F_{term} \mid F_{func} \\
F_{term} & ::= x \mid y \mid z \mid n_x \mid n_y \mid n_z \mid diff \mid ephem \\
F_{func} & ::= sin(F) \mid cos(F) \mid F + F \mid F - F \mid F * F \mid avg(F, F) \mid lum(RGB) \\
& \quad \mid max(F, F) \mid min(F, F) \mid if(F, F, F) \mid noise(RGB)
\end{aligned}$$

Fig. 4. Texture Language Definition

<u>Parameter</u>	<u>Value</u>
Population size	1000
Generations	200 (sometimes more)
Runs/experiment	10
Initialization	ramped half&half
Initial ramped tree depth	5 to 10
Max. tree depth	17
Crossover rate	0.9
Mutation rate	0.1
Selection scheme	tournament (size 5)

Table 1
Genetic Programming Parameters

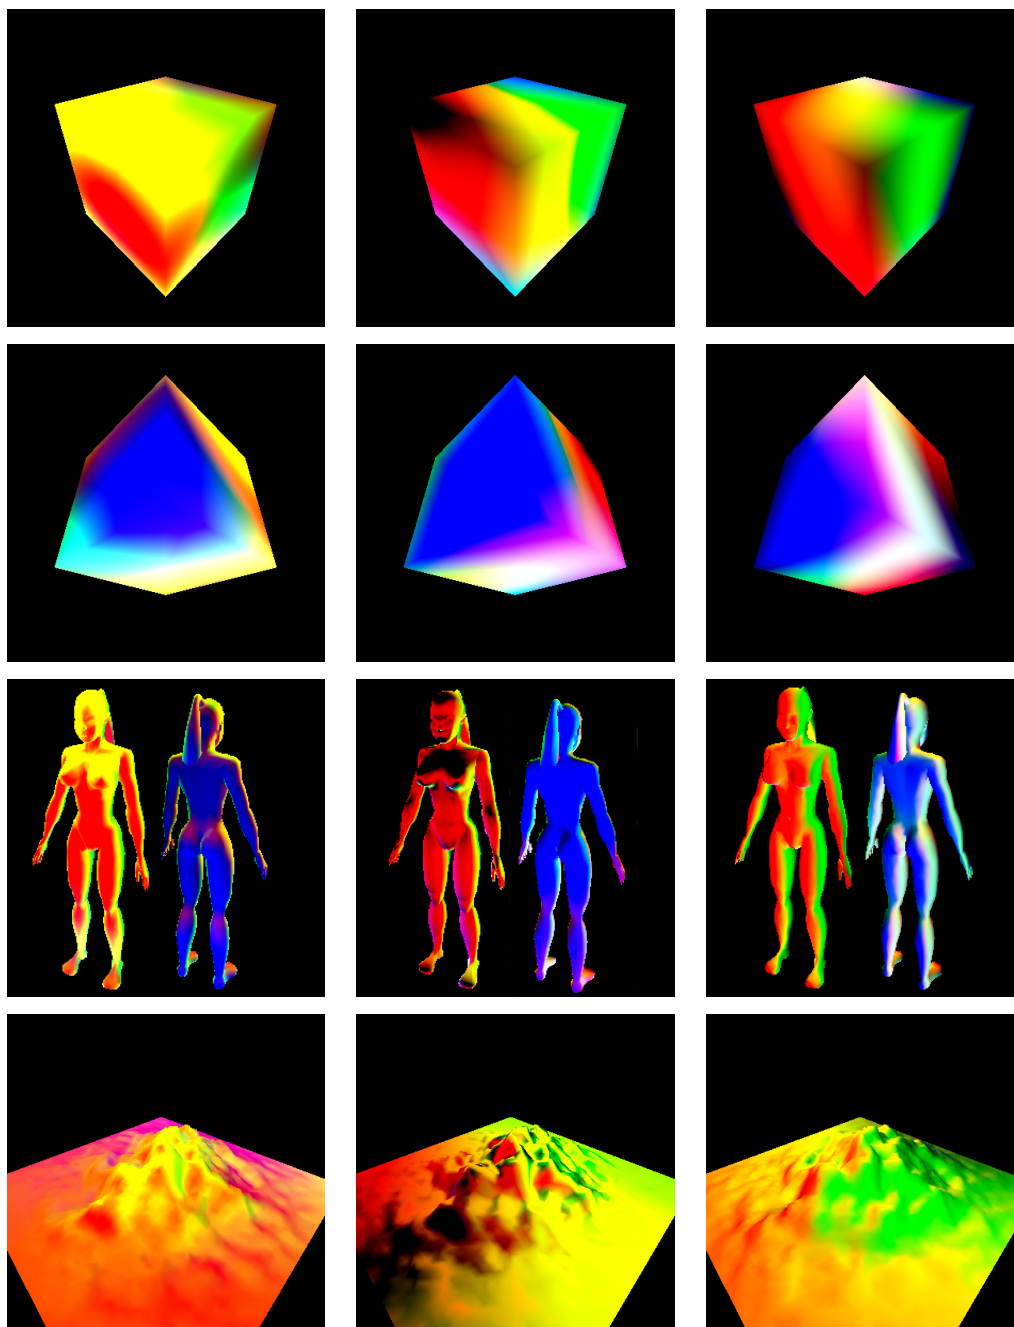


Fig. 5. Primary colour cube. Each column is a separate solution. The rows show training front view, training back view, woman front and back views, and terrain.

```

(rgb (+ (cos (+ (avg Z NZ) (- (* (- (cos (if (cos (- Z Y)) 0.28686 Y)) (min 0.36341 0.13969)) (min (+ (- Z NZ)
(* NX X)) (- Z NX))) NX))) (* (- (cos (if (- (cos (if (min 0.36341 0.13969) (- (min 0.36341 0.13969) (avg Z NZ))
(if (min NY 0.13969) (min (+ (avg Z NZ) (lum NXYZ)) (+ Z (- (min NY 0.13969) NZ))) Y))) 0.41939) (+ NZ
Z) Y)) (- Z NX)) (+ (* (avg (+ (- Z NX) (+ NY X)) (avg (+ (+ NZ (+ NY X)) (avg (+ (+ NZ NX) (+ NY X))
(- (cos (avg (+ (+ Z (* NX X)) X) NX)) (- (cos (min NY 0.13969)) (- Z NX)))))) (- (cos (+ (avg Z NZ) (+ Y (-
Z NX)))) (min 0.36341 0.13969)))) (- Z NZ)) (- NZ NX))) (cos (+ (min (avg (+ (min (+ Z (min (+ (avg Z NZ)
(lum NXYZ)) (cos (+ (* (avg (+ (+ NY X) (avg Z NZ)) (- Z NZ)) (- (cos NZ) NY)) (- NZ NX)))))) (avg (+ (+ Z
(* NX X)) X) NX)) (+ NY X)) (- (cos (min (avg (+ NY (* NX X)) (avg (+ NY X)) (- (cos (avg (+ NY (+ (+ Z
(* NX X)) X)) (- (cos (- Z Y)) (min 0.36341 0.13969)))) (- (cos (min 0.36341 0.13969)) (min 0.36341 0.13969))))
(* NX X)) (min 0.36341 0.13969))) (min (avg (+ (+ NZ NX) (+ NY X)) (- (cos (min (avg (+ (+ NZ NX) (-
(- (cos NZ) NY) 0.41939)) (- (cos (avg (- Z NX) NZ)) (- (cos (if (lum NXYZ) 0.28686 Y)) (- Z NX)))) (- (*
(- (cos (avg (- NZ NZ) NZ)) (- (cos (min 0.36341 0.13969)) (min 0.36341 0.13969)))) (- Z NZ)) NX) NZ))) (min
0.36341 0.13969))) (* (avg (+ (* NX X) X) (avg (+ (+ Z (* NX X)) X) NX)) X)) (+ Z (- (+ NZ (+ Z (* NX
X)) NX))) (* (- (avg (- (cos (+ (* (avg (+ (+ NZ NX) (+ NY X)) (- (cos (- Z Y)) (min 0.36341 0.13969)))) (-
(- NZ NX) NZ)) (- NZ NX)) (+ NY X)) (sin X)) (min (+ (avg Z NZ) (lum NXYZ)) (cos (cos (- (cos (cos (+ (*
(avg (+ (+ NY X) (avg Z NZ)) (- Z NZ)) (+ Z (* (- Z NX) X)) (- NZ (+ Z (min (+ (avg Z NZ) (lum NXYZ)) (-
(cos (* NX X)) (min 0.36341 0.13969)))))) (cos (- (+ (+ NZ NZ) (avg (+ (+ NZ NX) (+ NY X)) (- (cos (avg
(+ (+ Z (* NX X)) X) NX)) (- (cos (min NY 0.13969)) (sin X)))) (min 0.36341 0.13969)))))) (cos (avg (- (+
(- Z NZ) (* NX X)) (+ (avg Z NZ) (lum NXYZ))) (avg (max (+ (avg Z NZ) (+ Z (- NZ NX)) (min (+ (* (avg
(+ (+ NZ NX) (+ NY X)) (- (cos (- Z Y)) (min 0.36341 0.13969)))) (- (- NZ NX) NZ)) (- NZ NX)) (- (cos NZ
NY))) (- (* (- (cos NZ) (- Z NX)) (min (sin X) (- Z NZ)) NX))))))

```

Fig. 6. Example solution




Fig. 7. Terrain texture. Each column is a separate solution.




Fig. 8. Woman's clothing training points.




Fig. 9. Woman's clothing texture. Four solutions.

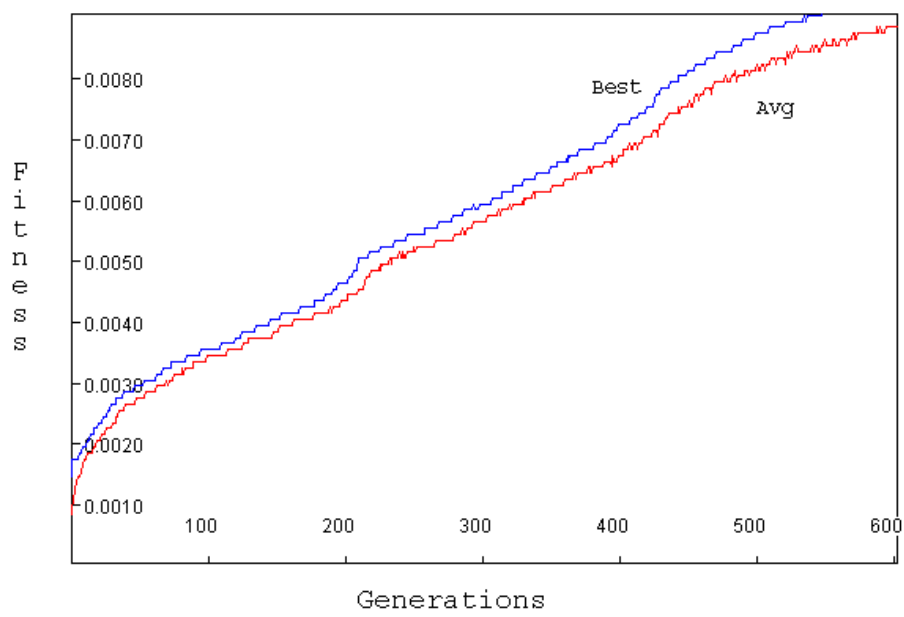


Fig. 10. Fitness progress for clothing texture run.




Fig. 11. Normal and gradient specialized texture. Each row is a solution.




Fig. 12. Swapping blue channel with gradient. Each row is a solution.




Fig. 13. Banded cube training points.




Fig. 14. Banded cube texture. Each row is a solution.