in Artificial Life VIII, *Standish, Abbass, Bedau (eds)(MIT Press) 2002. pp 157–163*

Aesthetic Selection of Developmental Art Forms

Dale Thomas
Artificial Intelligence Laboratory
Department of Information Technology
University of Zürich
Winterthurerstrasse 190
CH-8057, Zürich, Switzerland
thomas@ifi.unizh.ch

Abstract

Aesthetic selection has been employed in a system such that users can guide the evolution of art forms. The attributes of these forms are defined implicitly by an artificial genome, requiring a growth period to generate the final structure. During this growth phase, artificial chemicals react and diffuse across the surface of the developing structure, causing growth and other structural changes. The forms are part of a small evolving population in which fitness of each individual is entirely defined by the user. Throughout evolution user can design objects interactively or simply to explore the range of possible forms of the system.

Introduction

Genetic Algorithms (GA) were originally designed as a search technique, inspired by evolution using natural selection (NS) to 'breed' good solutions (Holland 1992). The concept is quite simple: from a population of information strings (genotypes) which define solutions (phenotypes) to a problem, a new population is created by breeding the solutions with a probability proportional to their fitness. This can lead to fitter populations, optimising solutions to the problem. Two elements of this method are very important in determining how well a GA performs: the mapping from genotype to phenotype and the fitness function. The fitness function defines the shape of the fitness landscape and the phenotype defines the position of an individual on it. In order for GAs to work, it is essential for the child solutions to be relatively similar to their parents. This requires that the fitness function should lead to a relatively smooth shape of landscape and that the mapping from genotype to phenotype must be such that a small change in genotype should generally lead to a small movement on the fitness landscape. Otherwise, evolution has nothing to work with because the fitness of the 'child' solutions, which are genetically similar to their parents, may be almost random, leading to an inefficient parallel random search. The advantages of GAs over other search techniques are that they are inherently parallel, simple to implement, do not require explicit knowledge of the problem, they can often find a good solution quickly and can search high-dimensional parameter space very well. The disadvantages are that they do not guarantee the optimal solution and that the performance very much depends on the shape of the 'fitness landscape', which is greatly influenced not only by the fitness function but also genetic encoding. In addition, the fitness landscapes of agent-based systems are further affected by co-evolution, changing environment/resources and sensory-motor interactions with the environment. These other factors can provide very strong evolutionary forces affecting the fitness of individuals in the system (Mitchell, Forrest, & Holland 1991). Also, the fitness will be greatly influenced by certain assumptions implicit in the implementation. Implicit information is a big problem when designing fitness functions and can lead to fragile solutions. For example, the AE system in (Bongard & Pfeifer 2001) had a simple fitness function in which the creatures were evolved to move toward a block in a simulated physical environment. The simulation was run for eight seconds and the fitness defined by the reciprocal of the final distance to the block. The distance measure was explicit in the fitness function but the timing was implicit. Evolution found it easier to make creatures move at a particular speed so as to be near the block at the end of the eight seconds evaluation period, rather than complex 'block-following' neural architecture. The fitness function should have been 'go toward the block' but the actual function was 'be here in eight seconds'. This simple example shows that even automated fitness functions can be highly biased due to implicit designing.

A central problem with GAs is that due to the algorithmic nature of fitness functions, they cannot be applied to problems with a subjective or qualitative aspect. Fitness functions such as size of a structure, speed of an agent, strength of a bridge, etc, are relatively easy to calculate but some things like 'looks like?' and 'appeals to me' are subjective choices and therefore impossible to encode. To overcome this limitation, another evolutionary selection method called Aesthetic (or Artificial) Selection (AS) is sometimes used. Many animal and plant species or have been subjected to this selection method since humans started to domesticate them. The fitness of individuals is based on human needs or desires. Some examples include dogs bred for 'beauty', hunting ability and affection, horses for speed and stamina and cows for milk yield. Also plants are selectively bred for yield and beauty. AS is a much simplified selection pressure since it concerns only very few aspects, whereas NS has to balance many aspects in order to keep the species alive. For this reason AS is faster at evolving specific traits. Also, NS has no goal and so just meanders whereas AS has a goal albeit biased and subjective. In Artificial Evolution (AE) systems this could potentially be a useful design tool for somewhat subjective problems.

In addition, since a human observer is controlling the selection, the fitness function can be flexible, have global information, and be adaptive and subjective which can sometimes lead to 'one step back for many steps forward'. For example, if the problem is to evolve agents for locomotion, a naive fitness function may be simply based on distance or speed of an individual. However, a human observer may see an individual that moves less far, but has an extra pair of legs or a particularly promising gait. An automated fitness function may not be able to detect this unless it was very complex (computationally intensive) but a human could easily see potential, even if the distance travelled was less.

Artificial Life (AL) systems have been used many times to produce interesting shapes and structures employing various methods including GAs. Usually GA methods include an explicit mapping from genotype to phenotype. This has the advantage of being intuitive and readable by the researchers but the main drawback is that the complexity of the phenotype is proportional to the length of the genotype. This is no real problem for small, manageable problems but as the size and complexity of artificial agents and structures increase, the disadvantage of this approach comes to bear with increasing force. Nature uses a different approach in which the complexity is not directly correlated to genome size. DNA does not directly encode the size, shape or function of the organism; rather it defines proteins and the organism's structure emerges through the interaction of these proteins and also the interaction with the environment.

The way that an organism grows is certainly one of the greatest mysteries of modern science. A seminal paper by Alan Turing (Turing 1952) described reaction-diffusion (RD) systems as a possible mechanism to explain some areas of morphogenesis. The information storage capacity of a system can be substantially less if there is a growth period exploiting self-organisation principles.

The work presented here combines the elements of artificial growth and artificial evolution using aesthetic selection into an education system for public use.

In the next section there is a discussion of related work and then the system is described starting with the interface, then a description of the physics and the chemistry and genetics are given in the subsequent two sections. Following that is a section on the graphics and then a general discussion on AS. Conclusions and future work are given in the final section.

Science et Cité

The system detailed here was exhibited in Zürich at the 'Science et Cité', a science fair held throughout Switzerland every year to promote public awareness of science. Many people tried out the system and were able to evolve something unique to their own aesthetic taste. Each user could save a picture of their form, which was later e-mailed to them in return for answering a few questions.

More information about this project can be found at: http://www.ifi.unizh.ch/ailab/people/thomas/

Related Work

There have been many interesting interactive art systems using artificial life. The first such system was Biomorphs (Dawkins 1986). This allowed the user to evolve binary trees where the genome consisted of nine integer values. To produce the phenotype, a recursive binary tree was rendered directly taking into account the genome which encoded things like recursion depth, branching angle, segment length, etc.

The user of the program was presented with a single biomorph surrounded by eight children with slightly mutated copies of the genotype. Once selected, a particular child would become the parent surrounded by eight of its children and the selection continued. This was simple asexual reproduction and even though it was simple, the author was able to direct the evolution to produce many desired shapes including the letters of his name.

Another similar system was an AS program for producing two-dimensional pictures from equations (Sims 1991). The equations were encoded as Lisp expressions and evolved using Genetic Programming (GP) (Koza 1992).

Sometimes some components of the fitness function are somewhat subjective while others may be quantifiable. Ventrella (1995) created a system that mixed automatic and aesthetic selection to evolve stylistic locomotion behaviour in simulated three-dimensional animats for computer animation. Some aspects of walking are objective, e.g. speed, distance/path travelled and efficiency whereas style is totally subjective. The fastest or most efficient method of moving may not necessarily be interesting from a subjective observer's point of view, which is of course very important in entertainment. In the system, the animator could breed a population for locomotion using standard AE techniques but also guide the evolution based on his or her whims. This was accomplished by being able to view a population and tweak the fitness after it had been calculated. This mixed the creative aspects of designers and animators with the automatic optimisation of evolution. Another advantage of this is that the user could help the system escape local minima if detected.

RD systems were originally designed to explain some areas of morphogenesis (Turing 1952). In these systems, there are interactions between diffusing chemicals, which are capable of a wide range of pattern formation. These systems can show stable spots, stripes, travelling waves, spirals, splitting structures and spatiotemporal chaos, and the patterns are very robust. Many such systems have been designed to model real chemical interactions or theoretical systems for use in biology, physics or mathematics. A lot of these systems have been used to explain some areas of biology including models of animal skin textures (Murray & Myerscough 1991), seashell patterning (Meinhardt 1994), nerve conduction (Fitzhugh 1962; Nagumo, Arimoto, & Yoshizawa 1962) and cellular aggregation (Vasiev, Hogeweg, & Panfilov 1994). In addition, RD systems have been used in computer graphics to model natural looking textures on an arbitrary, though static mesh (Turk 1991). Evolution has been used together with RD systems (Takai, Takai, & Nakamori 1998) but only in the optimisation of parameters for a specified system. The actual equations and parameters of RD systems have never been explored in an evolutionary context, mainly due to the requirement of positive and negative feedback in just the right amount.

The Interface

The interface to the system is similar to most other AS systems (Dawkins 1986; Sims 1991; Todd & Latham 1992; Ventrella 1994). It involves having a number of individuals on screen together (sixteen in this work) of which the user can select one to mutate, or two to breed. This then creates a new generation composed of individuals similar to the ones selected and the process repeats. This allows the user to have direct control over the evolution by exercising his or her own fitness value. Since this system was designed for public use using a touch screen, the interface needed to be as user-friendly and intuitive as possible. There are two views of the forms available to the user, one view showing all sixteen individuals and the other showing a close-up of just one. Other actions possible are to regrow an individual or to save its image. Regrowing an individual is included because the way it grows is an important part of the aesthetic appeal and since the system is completely deterministic the final form would be exactly the same.

Physics

A very important aspect of development is the action of physical forces from the environment, (such as a womb or egg shell) and from the material properties of the internal structure. For example, when a sheet of cells invaginates, it does so also because of physical buckling, not just as a result of chemical properties. Chemical and physical processes are two highly important and interdependent aspects of morphogenesis.

Chemical Number	Function	Affects
1	Growth Factor	Springs
2	Active Springs	Springs
3	Hair	Faces
4	Horns	Faces
5	Lights	Faces
6	Colour Component (red)	Vertices
7	Colour Component (green)	Vertices
8	Colour Component (blue)	Vertices

Table 1: The effect of the structural chemicals

The forms in this work are defined in a simulated three-dimensional environment as a surface mesh structure consisting of springs. These springs are generally passive and exert a force on the two connected vertices proportional to the deviation from the spring's rest length. Some springs can be active: this simply means that their rest length varies in proportion to a sine wave. The frequency and amplitude of this change is constant and equal for all active springs but the phase is determined by the time of creation for each spring. This allows for a 'rippling' effect over the form; for example, a bulge would travel down a long 'tail' structure since the time of creation increases toward the tip. This results in a very organic looking motion.

Initially, all forms begin as a standard, well-defined 'egg'. This is just a small cube with a chemical gradient across the eight verticies. As they grow, if a face or edge becomes too large the surface can split and increase its complexity in that area. Figure 1 shows a form growing from an initial egg. The normal of the surface at each vertex is important for two of reasons. Firstly, it helps to define the surface, which will also be discussed further in the discussion section, and secondly, all vertices are subjected to a small force along their normal which acts as a kind of internal pressure, keeping the structure from folding in on itself. Another force exists to keep the faces somewhat rectangular. This is just a small force on each corner of a face across the diagonal with the magnitude proportional to the difference of the diagonal lengths. These forces also help to keep the structure solid and inflated. All masses are defined as 1 and so the acceleration on a vertex is equal to the sum of the forces and the system was solved using simple Euler integration. A lot of damping was included for two reasons: to avoid instabilities when a vertex was connected to many springs, and also to give the motion a viscous liquid look for aesthetic appeal.

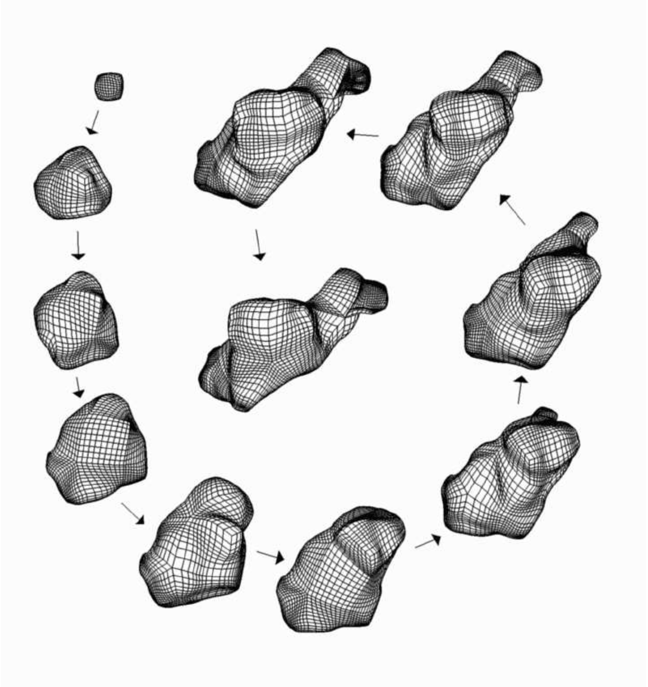

Figure 1: The growth of an individual.

Chemistry and Genetics

During development, the effect of physical forces is only one side of the story. The other side is the interactions between proteins and the differentiation of cells. Just as in biology, the genotype in this work does not explicitly encode for structure and detail; rather it encodes a description of the interactions between a number of abstract chemicals, N. Every vertex in the growing forms contains an N-dimensional vector of chemical concentrations and as the growth progresses the chemicals react locally with one another and diffuse through the springs with rates also defined by the genotype. The structure of the genotype is divided into two parts: the first N values are the diffusion rates of the respective chemicals and the rest is divided up into a number of genes, each consisting of five values. Each gene represents an interaction and these five numbers correspond to the FROM chemical, the TO chemical, the interaction weight and the upper and lower thresholds of the interaction. Each part of the genotype is represented as a real value in the range $[0, 1]$. These values are mapped onto the diffusion rates in the range $[0, 0.1]$ and the weights in the range of $[-0.1, 0.1]$. This creates an interaction network for the chemicals and as growth occurs, the reaction diffusion system affects the characteristics of the form. The interaction proceeds as follows: if the FROM chemical is between the limits of the interaction, the TO chemical is modified by the concentration of the FROM chemical multiplied by the weight. One main problem with evolving RD systems is that the systems need both positive and negative feedback in just the right amounts, otherwise the chemicals just decay to zero or saturate. This is very difficult to achieve and so the fitness landscape would be almost random. To solve this, each chemical was normalised over all vertices at every time-step. This is not biologically plausible but it did make the system stable. Also, this method is a very abstract model of transcriptional regulation in that one chemical can promote or inhibit the production of another. There are no other types of interactions such as enzymatic regulation.

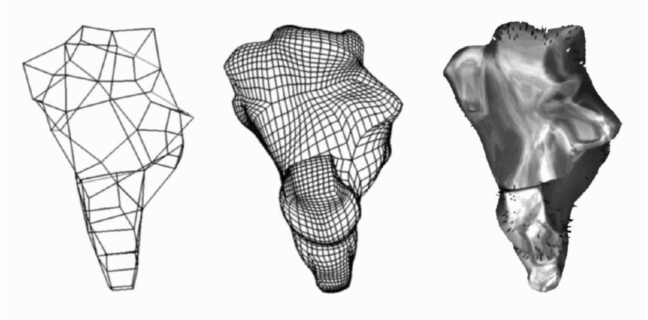

Figure 2: The three stages of object construction. On the left, the geometric wire frame spring structure is 'grown' from an initial cube. In the centre, the wire frame structure is used to generate a curved surface defined by the positions and normals of the vertices and finally this is rendered to the screen with many graphical effects.

Usually there are fifteen chemicals ($N = 15$), and the first seven are both structural and regulatory which means that they affect the physical form and also other chemicals while the rest are just regulatory. The functions of the structural chemicals are given in Table 1.

Since each form is uniquely defined by a string of numbers (artificial DNA), the standard genetic operators can be applied, i.e. mutation and crossover. In this work a simple one-point crossover was used and a variable mutation rate was used from 0.001 for the first child to 0.015 the fifteenth.

Graphics

The individuals in this work were evolved using AS so their appearance is obviously very important. Also, the motion and growth are aspects of the aesthetic appeal, which helps a user to decide which individual to choose and since the system is an entertainment package for public use the graphics were a major part. The graphics were implemented with OpenGL, a cross-platform graphics library with hardware support for fast 3D rendering.

In order to make interesting and unusual 'art' shapes, the overall look of the system was given an organic feel. The wire-frame mesh of the individuals was rendered with curved surfaces. Every face of the object was defined by four vertices and displayed as a curved surface by taking the positions of the vertices and their normals and calculating the positions of intermediate, internal vertices lying on the curved surface, which could then be rendered as much smaller quadrangles. These quadrangles were then rendered to the screen with appropriate lighting and texturing. Each vertex had a particular colour defined by the concentrations of three chemicals representing the red, green and blue components. A copy of the rotating background clouds was environment mapped onto the surface and the colours interpolated across it (environment mapping means displaying an object as though it was reflective). This gave the objects a slightly coloured chrome look. The resolution of this face division was dynamic and was controlled by the frame-rate of the system. If the objects became very complex then the system would begin to slow down and lower the resolution of the surfaces to keep the frame-rate relatively constant. Figure 2 shows the three stages of object construction. Beginning with the wire-frame spring structure, a curved surface mesh is generated and then rendered with texture and environment mapping, lighting and surface elements.

In order to increase the organic look of the system, natural looking elements like hair and horns were included to give subjective appearances of cute, disgusting, etc. These surface elements had no function other than aesthetic appeal and particular chemicals controlled the number of each element per face.

The hairs (or bristles) were rendered as simple black lines attached to the face and extending in the direction of the normal. The horns were similarly attached but rendered as a texture mapped pyramid. Also, orbiting lights were included and rendered as two quadrangle glare textures on billboards (i.e. they always faced the camera), rotating at different speeds and blended together. Each light was anchored to a fixed point on the surface as well but the distance varied along the normal as a sine wave with random phase. This gave a shimmering point cloud around the forms (depending on the chemical concentrations). As the form moved, (wobbling, rippling, growing, splitting), the normals were recalculated and so the elements move and wave with the form's motion. This gave a very unique look and helped with the aesthetics of the system.

If a face or edge became too long, it would split to extend the surface. Figure 3 shows the two splitting methods used. Since all faces are rendered as curved surfaces, the resulting faces after the splitting needed to have exactly four edges.

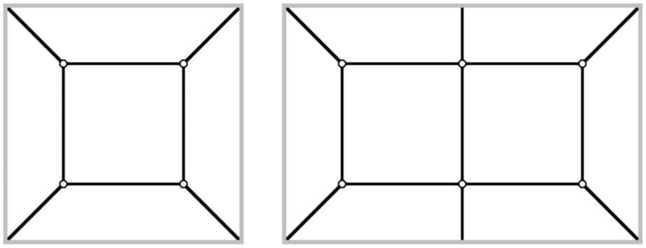

Figure 3: If a face becomes too large it will split into five new faces (left). If an edge becomes too long it will split into three new edges and also split the connected faces (right). All faces must be quadrangles. The grey lines are the original faces, the black lines show the new springs and the small circles show the newly created vertices.

The forms here can have active springs spanning across the interior of an individual. When they are created, they are created independent of the structural springs. Then if the structural spring in the same place splits due to growth, the active spring spans across the inside of the form.

Some sound was included for a more complete system. This just consisted of background noises: wind, bubbling, bird chirping and also a Taiko (Japanese drumming) song. The sounds did not relate to the forms in any way, except for an organic bubbling/crunching sound for when the user interacted with the forms.

Discussion

The main advantages of using a developmental mapping from genotype to phenotype is that the solution can adapt to environment during growth, and that the complexity of phenotype is not constrained to the length of genotype. Unfortunately, in artificial algorithms without the complex growth feedback checking available to biological systems, small changes in genotype can sometimes lead to large movements across the fitness landscape. However, by using an AS method this problem can be reduced due to a subjective, adaptive and sympathetic fitness function. Another advantage of this type of design method is that any fitness function can be implemented at any time. Here is a system where the fitness is based on the user's whims, it is unquantifiable and highly subjective but this can be a good thing in many ways. Also, since this is 'Art by Criticism' the user (or designer) does not need to have great artistic ability; he or she just needs some aesthetic sense. This may be useful for some areas of engineering or design where the computer could reduce the amount of technical ability needed by the user.

By 'playing around' with the system one can get a good intuition of the shape of phenotypic space, to see what is possible and what is easy to produce. The motivation can be to 'break' the system or just to explore it. The user can try to find the weak points, limitations and range of the system. For example, the work here uses an implicit mapping and so the phenotypic space is not at all intuitive.

AS is needed in some situations so that a simple, artificial system can find solutions that not only optimise their fitness criteria but also have aesthetic or natural appeal. Biological structures and behaviours evolved as a response to many different interactions with the environment and other agents. The world is a rich tapestry of dynamically changing variables due to physical forces and other organisms competing for resources and coevolving. In an artificial system it is either impossible, inefficient or undesirable to include such complexity just to be able to evolve natural-looking structures or behaviours. However, without such complexity it may be too much to ask that the resulting solution looks natural.

A general problem with developmental algorithms is that of growth termination. In biology there are many ways for an organism to stop growing after reaching maturity. For example, terminal differentiation of cells, which disables their ability to divide, the activity range of chemicals, environmental constraints, etc. In this work this issue was not addressed; rather the growth will terminate after a particular time. After this growth phase, the chemicals can still react and diffuse but only affecting the colours, so no more growth can occur. This would be a good area of future work, as it seems to be a very important aspect of development but as yet not included in artificial growth algorithms.

There are two potential problems with automated fitness functions, which can hinder the performance of a GA. Firstly, implicit information can lead to solutions which are not what the designer wanted, simply because the fitness function was not specific enough, due to the designer, expecting evolution to design as he does, and secondly, the designer can over-specify the function, restricting evolution to the his way of thinking and inhibiting evolutions inherent innovation. These are both problems of bias and as a result, many fitness functions are not truly objective. AS is subjective but it is also adaptive, so that sometimes, implicit information can be recognised and dealt with. Presumably, a mixture of AS and NS such as described by (Ventrella 1994) can overcome many problems with both approaches.

An additional feature of the program was the ability to ‘touch’ an individual. Touching the screen applied some perturbation forces to the vertices, which would initiate a few seconds of damped rippling. The main reason for this was to better connect the user to the forms. The idea of the system was to create ‘art forms’ but the use of evolution meant that the user is not in contact with the individual forms, rather just the whole population. The result of this is that the process of design has moved away from hands-on sculpting. However, most people like to be able to interact with the forms directly.

Conclusions and Future Work

Since this project was primarily designed as an interactive art project, the results are the images created by the users activities.

There are many possible directions for future research. One possible direction is the inclusion of an automatic selection system for a particular problem like locomotion (Sims 1994; Ventrella 1994; Bongard & Pfeifer 2001). Implementing gravity and collisions with the environment/self/other agents could enhance the physical environment. In all of these systems so far, the agents under evolution have been rigid body hierarchies. However, the material properties of an agent play an important role in its behaviour (Hara & Pfeifer 2000), and this system could be used to study this aspect in an evolutionary context. One point which would need to be addressed, is the need for some internal structure. In this work, the forms can have internal active springs but this is not enough to evolve interesting locomotive behaviour. What is needed is a rigid body skeleton covered with a soft ‘skin’. Then for these agents the control, morphology and material parameters could be evolved in synchrony.

Acknowledgements

I would like to thank Rolf Pfeifer and Josh Bongard for interesting discussions on various topics pertaining to this work. Also thank you to Daniel Bisig for very helpful suggestions and advice, and also for providing the sounds used in the system.

References

Bongard, J. C., and Pfeifer, R. 2001. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In *Proceedings of The Genetic and Evolutionary Computation Conference, GECCO-2001*, 829–836. San Francisco, CA: Morgan Kaufmann.

Brooks, R., and Maes, P., eds. 1994. *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems.* Cambridge, Mass.: MIT Press.

Dawkins, R. 1986. *The Blind Watchmaker.* Harlow, UK: Longman Scientific and Technical.

Fitzhugh, R. 1962. Impulse and physiological states in models of nerve membrane. *Biophysics J.* 1:445–466.

Hara, F., and Pfeifer, R. 2000. On the relation among morphology, material, and control in morphofunctonal machines. In *From animals to animats 6. Proceedings of the 6th Int. Conf. on Simulation of Adaptive Behavior*, 33–42. Cambridge, Mass.: MIT Press.

Holland, J. 1992. *Adaptation in Natural and Artificial Systems.* Cambridge, MA: MIT Press, 2nd edition.

Koza, J. R. 1992. *Genetic Programming - on the programming of computers by means of natural selection.*

MIT Press.

Meinhardt, H. 1994. Biological pattern-formation — new observations provide support for theoretical predictions. *Bioessays* 627–632.

Mitchell, M.; Forrest, S.; and Holland, J. 1991. The royal road for genetic algorithms: Fitness landscapes and GA performance. In *Proceedings of the First European Conference on Artificial Life*, 245–254. Cambridge, MA: MIT Press.

Murray, J. D., and Myerscough, M. R. 1991. Pigmentation pattern formation on snakes. *J. Theor. Biol.* 149:339–360.

Nagumo, J. S.; Arimoto, S.; and Yoshizawa, S. 1962. An active pulse transmission line simulating nerve axon. *Proc. IRE* 50:2061–2071.

Sims, K. 1991. Artificial evolution for computer graphics. In *Computer Graphics (Siggraph '91 proceedings)*, 319–328.

Sims, K. 1994. Evolving 3D morphology and behavior by competition. In Brooks and Maes (1994), 28–39.

Takai, Y.; Takai, N. K.; and Nakamori, K. J. 1998. Exploration of the reaction-diffusion textures. In *Proc. Int. Conf. on Modelling and Simulation.*

Todd, S., and Latham, W. 1992. *Evolutionary Art and Computers.* London, San Diego CA: Academic Press.

Turing, A. 1952. The chemical basis of morphogenesis. *Phil Trans R Soc B* 237:37–72.

Turk, G. 1991. Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics* 25(4):289–298.

Vasiev, B.; Hogeweg, P.; and Panfilov, A. 1994. Simulation of dictyostelium discoideum aggregation via reaction-diffusion model. *Physical Review Letters* 73:31733176.

Ventrella, J. 1994. Disney meets Darwin. In Brooks and Maes (1994), 436–441.