

A Gene Regulatory Network for Cell Differentiation in *Caenorhabditis elegans*

Nicholas Geard¹ and Janet Wiles^{1,2}

¹ School of Information Technology and Electrical Engineering

² School of Psychology

The University of Queensland,
St Lucia, Queensland 4072, Australia
{nic, j.wiles}@itee.uq.edu.au

Abstract. Biological development is a remarkably complex process. A single cell, in an appropriate environment, contains enough information to produce a wide variety of specialised cell types, whose spatial and temporal dynamics interact to form intricately detailed patterns and behaviour.

Much of the complexity of a developing system lies in the dynamics of gene regulation that occur within each cell. We used a simple recurrent network to model the process of gene regulation and evolved systems that were able to generate the first four cell divisions of the *C. elegans* cell lineage tree with a high degree of accuracy.

1 Introduction

The development of a single, fertilised egg cell into an elaborate, multicellular organism is one of the most complex processes in biology [36]. Cells are able to achieve and maintain stable patterns of gene activation over a large number of cell divisions while continuously integrating information from their environment and signals from other cells. Their activities are carefully coordinated through both space and time to produce a correctly formed organism, all without any form of centralised control.

Some of the most interesting areas of investigation within developmental biology include how cells differentiate over time, how body plan patterning and structure emerge from the internal dynamics and external interactions of groups of cells, how the information required to guide this process is encoded in an organism's genome and how developmental processes change over evolutionary time.

Computer modelling and simulation have the potential to contribute to these investigations. They can enable the synthesis and analysis of vast amounts of experimental data, such as that produced by genome sequencing projects and gene expression profiling. Computational models also allow hypotheses about evolutionary processes and emergent properties, such as robustness and evolvability, to be generated and explored.

One field that has made extensive use of the concept of development is artificial life, which recognises it as a powerful and highly evolvable means of encoding design solutions [3]. This view has led many researchers in artificial life to approach biology primarily as a source of novel ideas capable of revolutionising engineering and other domains [25]. However, we believe that artificial life models also have the potential to make a contribution to biology (see also [21]). One issue that must continually be addressed is that of bridging the gap between computational models of biological processes and specific biological tasks in a way that makes artificial life models accessible to a wider scientific audience.

In this study, we introduce a simplified gene regulatory network model based on a standard recurrent network control architecture and demonstrate the ability of our model to generate the beginnings of the cell lineage tree of the nematode worm, *Caenorhabditis elegans*.

The next section describes cell differentiation and the way in which patterns of gene activation determine a cell's fate. The early embryogenic steps of *C. elegans* are then described in more detail. Previous computational approaches to modelling gene regulation and development are briefly reviewed before our own recurrent network model is presented. The ability of our model to evolve a network capable of generating a set of differentiated cells is demonstrated and an analysis of its performance with respect to accuracy and search time is reported.

2 Mechanisms of Cell Differentiation

Cell differentiation is the process by which cells undergo physical and chemical changes that result in them becoming structurally and functionally distinct. From an initial, undifferentiated egg cell, each successive cell division results in a new generation of cells whose final fate is more determined. Upon reaching its terminal, fully differentiated state, a cell will function as, for example, a nerve, muscle or blood cell.

The primary feature that determines the function of a fully differentiated cell is the proteins it contains [1]. Similarly, the most important property characterising a developing cell is its pattern of gene activity. In the early stages of embryo development, the few cells produced are unlikely to be visibly distinct from one another. Already, however, the differences in the patterns of gene activity of these precursor cells will determine the role they and their daughters are to play in the fully developed organism.

When a cell divides, its pattern of active and inactive genes is passed on to its daughter cells via a number of different mechanisms [36]. The sections below describe how eukaryotic genes are regulated, how patterns of activity are maintained through a cell lineage and how these patterns can change as each cell's fate becomes more determined. Several other mechanisms that play a part in the differentiation of cells are also described.

2.1 The Control of Cell Differentiation by Gene Regulation

The process of gene expression begins when an RNA polymerase molecule binds to the start site of a gene, unwinds a section of DNA and uses one of the strands as a template to transcribe messenger RNA (mRNA) molecules. mRNA molecules are transported outside of the cell nucleus to the cytoplasm, where they are translated into proteins by ribosomes. Proteins can either be functional, enabling a cell to carry out the chemical processes required by its role as, for example, a nerve or blood cell, or they can re-enter the nucleus to regulate the expression of other genes.

Regulatory proteins interact with the promoter and control regions of a gene to either activate or repress the transcription of that gene. Some regulatory proteins, known as general transcription factors, are required for any transcription to occur at all. Others play a role as activators, binding to enhancer sites located upstream or downstream of the gene to facilitate transcription. Yet another type of regulatory protein acts as a repressor, either blocking the activity of activator proteins, or preventing the binding of RNA polymerase to a gene start site from occurring at all [16].

When a cell divides, the set of regulatory proteins that determine its pattern of gene activation are divided between the daughter cells, so each will generally have a similar pattern of gene expression to its parent. Patterns of gene activity can also be maintained across cell divisions by heritable chemical and structural modifications to the genome. DNA methylation, a chemical modification to DNA, can inactivate individual genes, and chromatin structure can be repackaged, rendering localised regions of the genome transcriptionally inactive.

While the maintenance of patterns of gene activation is an important part of development, for cell differentiation to occur, it is necessary that some differences arise in the activity of newly formed cells [36]. There are several ways that this can occur. As well as being regulated by a cell's own internal chemistry, genes can also be switched on and off by external signals originating from other cells or from the environment. In general, these signalling molecules bind to receptors found on the cell surface, and the signal is transmitted to the nucleus via a series of chemical events called a signal transduction pathway. Asymmetric cell division is yet another way that differences between cells may arise. When a parent cell divides, the two daughter cells may not necessarily be identical. Irrespective of any influence from the cellular environment, the concentration of regulatory and other proteins in the parent cell's cytoplasm may be unevenly distributed between daughters. As well as resulting in cells of different sizes, this inequality also has the potential to affect the pattern of gene activation of each daughter in a different manner.

2.2 Other Mechanisms of Controlling Cell Differentiation

In addition to gene regulation, other mechanisms involved in cell differentiation include induction and the use of positional information.

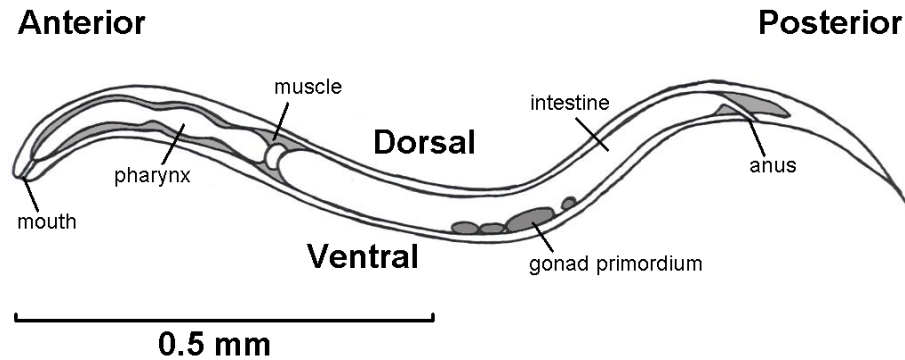


Fig. 1. *Caenorhabditis elegans* as a newly hatched larva. The pharynx is used to pump food from the mouth to the intestine. The nervous system consists primarily of a number of sense organs around the head. Four bands of muscle running the length of the organism allow it to propel itself by alternatively flexing and relaxing. The gonad primordium will continue to develop into the reproductive system. The precursor cells specifying each of these different anatomical features are, for the most part, formed only a few cell divisions after the fertilisation of an egg cell. Redrawn from [36].

Inductive interactions between adjacent groups of cells play an important role in cell differentiation by altering the fate of the target cells. The signals involved may either act locally, via direct contact between cells, or be transmitted over a greater distance via diffusible molecules. A controlled series of induction steps is one of the approaches utilised by development to specify the body plan of an organism. A second mechanism responsible for pattern formation is the use of positional information. Several theories have been proposed for how cells may determine their position. One possible method would be by tracking a chemical whose concentration varies in a regular fashion over a region of cells. Another approach would be to use local signals from surrounding cells to determine relative positional information.

3 The Early Embryogeny of *C. elegans*

C. elegans is a small (approximately 1mm as an adult) worm that lives in rotting vegetation around the world and survives on a diet of bacteria (Figure 1). It is an important model organism for developmental biologists for a number of reasons. It has a relatively small number of cells (less than a thousand as a fully developed adult) with an invariant cell lineage, and a short life-cycle, making it highly suitable for laboratory experimentation. Furthermore, it is transparent, allowing observation of the formation and movement of cells throughout development. Most significantly, it is one of the simplest existing organisms that shares many of the same biological characteristics as humans.

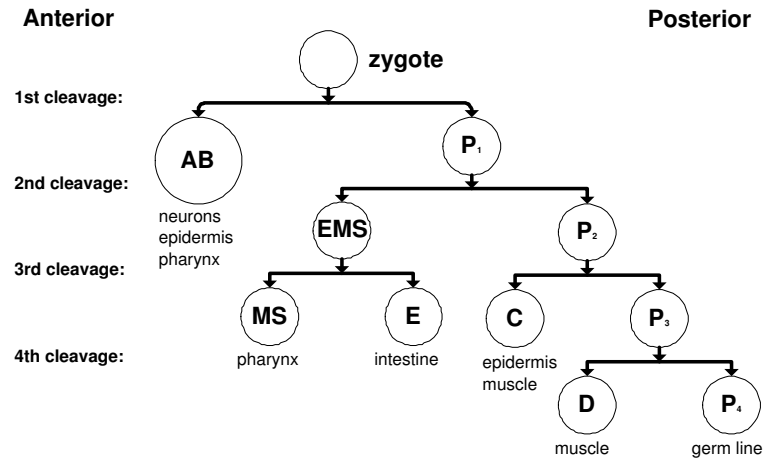


Fig. 2. The cell lineage tree of very early *C. elegans* embryogeny. Each precursor cell cleavage results in the production of one somatic cell and a further precursor cell. The final precursor cell, P₄, gives rise to the germ line. See text for further details. Redrawn from [36].

The first complete observation of *C. elegans* embryogenesis was carried out by Sulston and colleagues in 1983 [28], resulting in a tree of the entire cell lineage. The initial cell divisions are invariant, with each precursor cell cleavage producing one somatic cell, which goes on to specify some region of the embryo, and a further precursor cell. After the fourth cleavage, the precursor cell forms the basis of the organism's germ line [36] (see Figure 2).

First cleavage: Before fertilisation, there is no asymmetry in the egg cell. The point of sperm entry determines the first, asymmetric cleavage of the zygote and defines the orientation of the anterior–posterior axis. The two daughter cells produced are the AB cell, which goes on to produce cells related to the epidermis, nervous system and muscle tissue, and another precursor cell, P₁.

Second cleavage: At this point, the P₁ cell divides to produce the EMS cell and a further precursor cell P₂. The relative positions of the cells at this point define the orientation of the dorsal–ventral axis.

Third cleavage At the third cleavage the EMS cell divides to form the MS and E cells. MS goes on to specify a portion of the pharynx while the E cell gives rise to the gut. The P₂ cell divides to produce the C cell, which specifies the epidermis and muscle, and P₃, another precursor cell.

Fourth cleavage Finally, P₃ divides to form the D cell, which produces muscle, and the P₄ cell, which gives rise to germ cells.

At this early stage, the major regions of the developing embryo have already been specified. The role that intra-cellular interactions play in the process of differentiation has been demonstrated by manipulating an embryo after the first cleavage to change the orientation of the AB cell [36]. Not only does this reverse the position of the daughters of the AB cell, but the position of the P₁ daughter cell, EMS, is also reversed relative to the AB cells, resulting in a reversal of the dorsal–ventral axis.

The primary mechanisms determining the fates of the first few cells in the *C. elegans* embryo are therefore the cell lineage, specified by the change in patterns of gene activation, and cell interactions.

4 Methods for Modelling Gene Regulation and Development

Recent years have seen a rapid increase in the use of modelling as a tool to understand biological processes. Approaches to modelling range from detailed, mathematical models tied closely to experimental data (see [6, 32] for recent reviews) to more abstract models from the fields of complex systems and artificial life that aim to develop insights about the high-level properties of gene regulation (e.g., [18, 26, 35, 5]).

The sections below focus on reviewing the network approach to modelling gene regulation and the various models that use gene regulation to control developmental processes.

4.1 The Network Approach to Gene Regulation

As described above, the regulation of gene expression is coordinated by interactions between multiple transcription factors that bind to a gene’s control region. Frequently, these transcription factors are also the products of gene transcription events, and hence they too are under regulatory control. The transcription of a given gene can therefore be defined as a combinatorial function of the levels of transcription of the other genes in the system, together with any environmental influences. Thus future states of the system can be predicted from the current pattern of gene expression.

More formally, a genetic system can be viewed as a network in which each node is a gene and each link defines a regulatory interaction. The activation of a node is updated based on the activations of each of the other nodes to which it is connected. In the simplest case, nodes are either on or off (i.e., either transcribed or not) and the condition for activation of a node is represented as a Boolean function of the input nodes. Depending on their size, connectivity and node update function, such systems are capable of displaying a variety of dynamic behaviours, ranging from ordered fixed-point and cyclic attractors to disordered or “chaotic” dynamics [18, 37, 8].

Other previous studies of gene regulation have used networks with continuous, rather than Boolean, activation [22, 20, 33, 34, 24]. These networks are sim-

ilar to a widely studied class of artificial neural network models known as recurrent neural networks [10]. Continuous activation networks capture several properties of gene regulatory networks not represented by Boolean networks. Firstly, genes may have a different effect depending on their level of expression that a simple on/off distinction does not capture. Secondly, genes may influence the transcription of different genes by differing amounts.

4.2 Modelling Development

Several studies have incorporated a developmental mapping between genotype and phenotype layers in a number of different domains including evolving neural networks [19, 23, 14], evolving autonomous agent and robot controllers and morphologies [17, 7, 9, 4] and other evolutionary design tasks [2, 31].

A number of artificial life researchers have also been interested in developing models of biological processes for the further insights that these models may provide about biology [11, 15, 12] (see [27] for a recent review).

An essential, but often challenging requirement for modelling any biological system is to determine an appropriate level of abstraction for the issues under investigation. At the “appropriate” level, properties that the model seeks to explain or investigate can be seen to emerge from interactions between well-understood agents at a lower level of description. The goal is to balance simplicity with biological plausibility. If a model is too simplistic, extensions to biology may be unconvincing, and findings will lack cogency. Conversely, overly elaborate models may include unnecessary details that obscure the essential dynamics of the process under investigation.

The core issue in our model of gene regulation and development is to understand how a network model of gene regulation can control the developmental sequence of a complex organism.

5 Methodology

5.1 The Recurrent Network Model

In our model, a genetic system was defined as a network of N interacting nodes (see Figure 3). The activation state of each node was a continuous variable in the range $[0, 1]$, where 0 represented a completely inactive gene and 1 a fully expressed gene. The network was updated synchronously in discrete steps, where each time step represents a single cell division.

In most organisms, the time-scale for a single cell cycle is measured in minutes or hours [1]. The time taken for an individual transcription event is considerably shorter, therefore a single cell cycle will consist of multiple transcription events. To capture the potential complexity of the interacting transcription factors, we have used a fully connected network. Thus an individual link in the network does not necessarily represent an actual physical interaction, but rather the degree of influence that the transcription of the source gene at time t has on

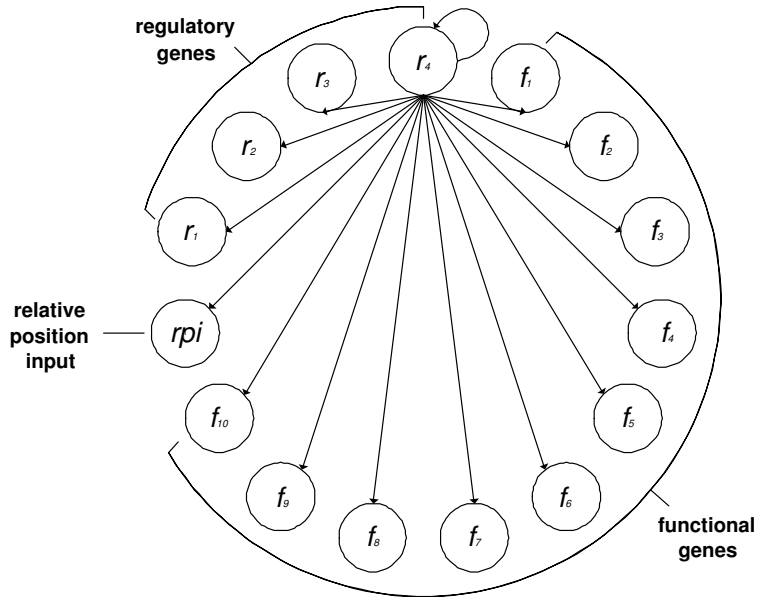


Fig. 3. The structure of the network. Gene regulation is modelled using a fully connected network of fifteen nodes. Section 5.1 describes the different types of gene in more detail. Note that only the regulatory links of the fourth regulatory gene are shown. The links of each of the other genes follow a similar pattern, but are omitted for clarity.

the transcription of the target gene at time $t + 1$. This influence is described by a weight on the link, which may be positive or negative, depending on whether the source gene is an activator or a repressor in that context. The inclusion of self-connections (i.e. from node i to node i) allows for the possibility of genes being self-regulatory. The inclusion of regulatory outputs from functional nodes accounts for any potential regulatory influence that their transcription may have, whether it be via direct interaction, competition for metabolic resources, or indirect feedback through the environment.

The state of the network was updated synchronously, with the activation of node i at time t , $a_i(t)$, given by

$$a_i(t) = \sigma\left(\sum_{j \in N} w_{ij} a_j(t-1) - \theta_i\right) \quad (1)$$

where w_{ij} is the strength of the interaction from node i to node j , θ_i is the activation threshold of node i , and $\sigma(\cdot)$ is the sigmoid function, given by

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

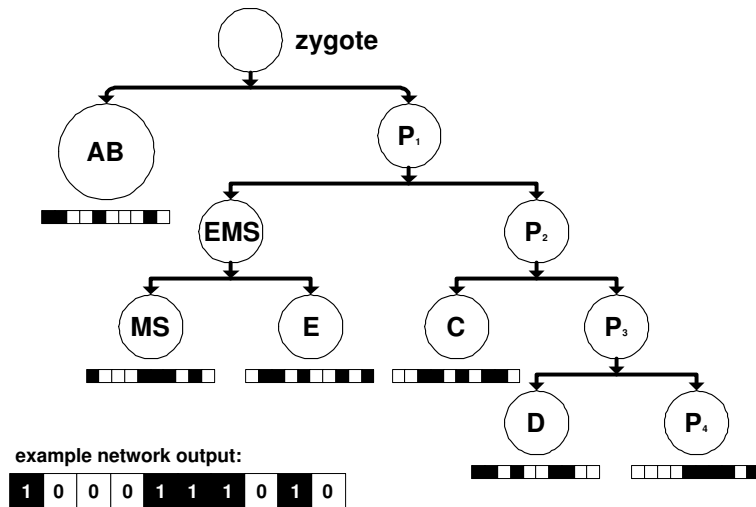


Fig. 4. The tree of target gene activation patterns. The bars under each terminal cell represent the pattern of gene activation that defines that cell's current state of differentiation. The enlarged example pattern illustrates the correspondence between the pattern and the target activation of the functional output nodes.

The network contained 15 nodes. One of the nodes was defined to be a relative position input. After a cell division, this node is set to 0 in the left daughter and 1 in the right daughter. This external input reflects the difference in intra-cellular and environmental signals to the cell resulting from their respective positions in the embryo. The next ten nodes were defined to be the functional outputs of the cell. These nodes represent a subset of genes whose pattern of activation specifies the current state of differentiation of the cell. The remaining four nodes in the network represent genes that play a regulatory role only. The model network contains considerably fewer nodes than the actual number of regulatory genes in *C. elegans*, however it was considered appropriate for the complexity of the task used in this study.

5.2 The Cell Differentiation Task

The task we required our system to perform was to generate a subset of the cell lineage tree of *C. elegans*, up to the point at which the major regions of the embryo are specified. At the beginning of each simulation run, we defined six random patterns of gene activation to represent the target pattern of gene activation for each terminal cell (see Figure 4).

A cell lineage tree was generated in the following way. Beginning with the original network, a single update was performed. Cell division then occurred, in which two copies of the network were created with identical weights and node

activations except for the relative position input node, which was set to 0 in the left daughter and 1 in the right daughter. This process was repeated for each non-terminal daughter cell until networks corresponding to cells AB, MS, E, C, D and P₄ were obtained. The activation of the ten functional genes in these cells was compared to the corresponding target pattern and the accuracy of this match was used as the basis for evaluating the network.

Two related error values were calculated based on the distance from the desired pattern. The first error value was the Sum Squared Error (SSE), which measured the difference between the target activation (either 1 or 0) and the continuous-valued activation of each node. The SSE was calculated by

$$SSE = \sum_{j=1}^6 \sum_{i=1}^{10} (p_i^j - f_i^j)^2 \quad (3)$$

where p_i^j is the activation of gene i in the target pattern j and f_i^j is the activation of the functional gene i in the network corresponding to the cell j . This error value was used as the basis for comparing two networks in the evolutionary trials described below.

A second error measure, the Number of Gene Errors (NGE) was defined as the number of incorrect gene activations in a pattern, where a correct gene activation was defined as being greater than 0.5 if the target activation was 1, and less than or equal to 0.5 if the target activation was 0. The NGE was calculated by

$$NGE = \sum_{j=1}^6 \sum_{i=1}^{10} \phi(p_i^j, f_i^j) \quad (4)$$

where

$$\phi(p_i, f_i) = \begin{cases} 0 & \text{if } ((p_i = 1) \wedge (f_i > 0.5)) \vee ((p_i = 0) \wedge (f_i \leq 0.5)) \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

A simple evolutionary search strategy called the 1+1 EA was used. Initially a single network was generated with uniformly distributed random weights in the range $[-1, 1]$. The error values for this individual were calculated as described above and stored. A mutant network was derived from this network by adding a uniformly distributed random value in the range $[-0.5, 0.5]$ to each weight. The error values for the modified individual were calculated and compared to that of the original individual. The individual with the lowest SSE was retained and used as the basis for the creation of a new mutant. This process repeated for 20,000 generations.

6 Results and Discussion

Twenty trials were run, each initialised with a different, randomly generated individual and a different, randomly generated set of target activation patterns.

All trials evolved a solution that was able to perform the cell differentiation task with above 90% accuracy within 20,000 generations (see Figure 5 for a time series of the SSE and NGE over a typical run). In 6 of the 20 trials, a solution containing only one incorrect gene (98.3% accuracy) was found.

Some variation was noted in the number of residual incorrect genes at the end of 20,000 generations (see Figure 6). When a good solution was found, it was typically found very quickly. The average time required to find a solution that was at least 95% correct (i.e., three or less residual errors) was only 3,308 generations (standard deviation = 2,304 generations).

We also examined progress toward the solution in a number of runs (see Figure 7). The evolutionary dynamic typically begins with an initial transient period, corresponding to the initial rapid decrease in error (see Figure 5 inset). This reduction in error is followed by a long period of relative stability as the search process fine tunes the network weights. In the example shown in Figure 7, a solution with only one error (in the ninth gene of the MS cell) was found after approximately 3,000 generations. Over the succeeding 17,000 generations, the SSE continued to decrease, but the residual error remained.

7 Conclusions and Further Work

The simulations demonstrate how a network model of genetic regulation can generate the gene expression patterns required for cell differentiation and development. By working with a problem taken directly from the domain of biology, we have taken steps to bridging the gap between computational models and real data.

As it is presented here, our recurrent network model and task definition represent a high level of abstraction from the biological systems under consideration. In current work, we are extending the approach described above to explore alternative network structures, and target patterns derived from biological data, with the ultimate aim of evolving a system capable of generating the entire *C. elegans* cell lineage tree. This more complex task has raised several interesting questions about the dynamics of gene regulation in both artificial and natural systems [30].

Furthermore, with the sequencing of the *C. elegans* genome complete [29], it should be possible to extend our knowledge of the embryogenic process to understand how developmental information is encoded in the genome, and hence how it interacts with evolutionary processes. In a complementary study, we have been exploring the use of computational models of genome sequences as a tool for the investigation of regulatory network structure and dynamics [13].

Acknowledgements

This work was supported by an APA to NG and an ARC grant to JW.

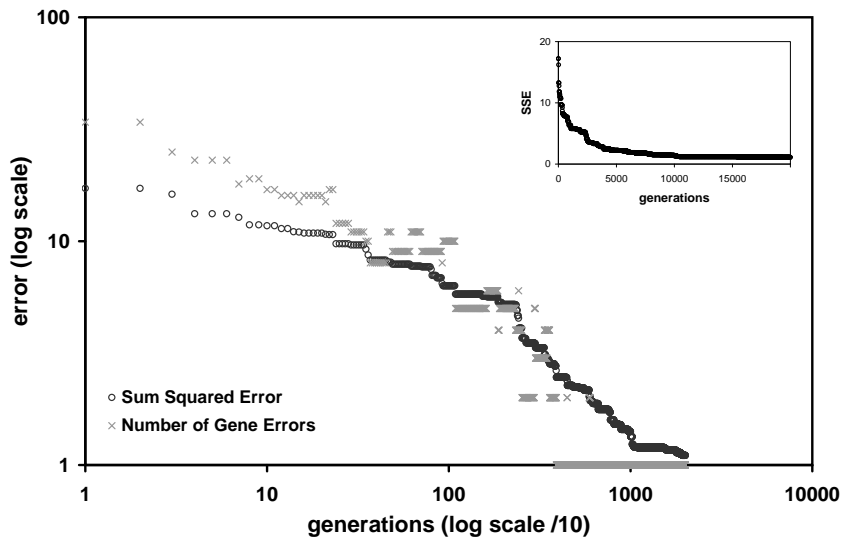


Fig. 5. Performance in a typical evolutionary run. The number of incorrect genes generated by the system rapidly reduces to 1 (by approximately generation 4,000) and further evolution results in the SSE continuing to decrease. SSE was used as the selection criterion and therefore decreases monotonically. NGE fluctuates slightly as better individuals under this measure are occasionally lost. The inset shows a linear plot of the SSE.




Fig. 6. A summary of the number of residual errors after 20,000 generations.

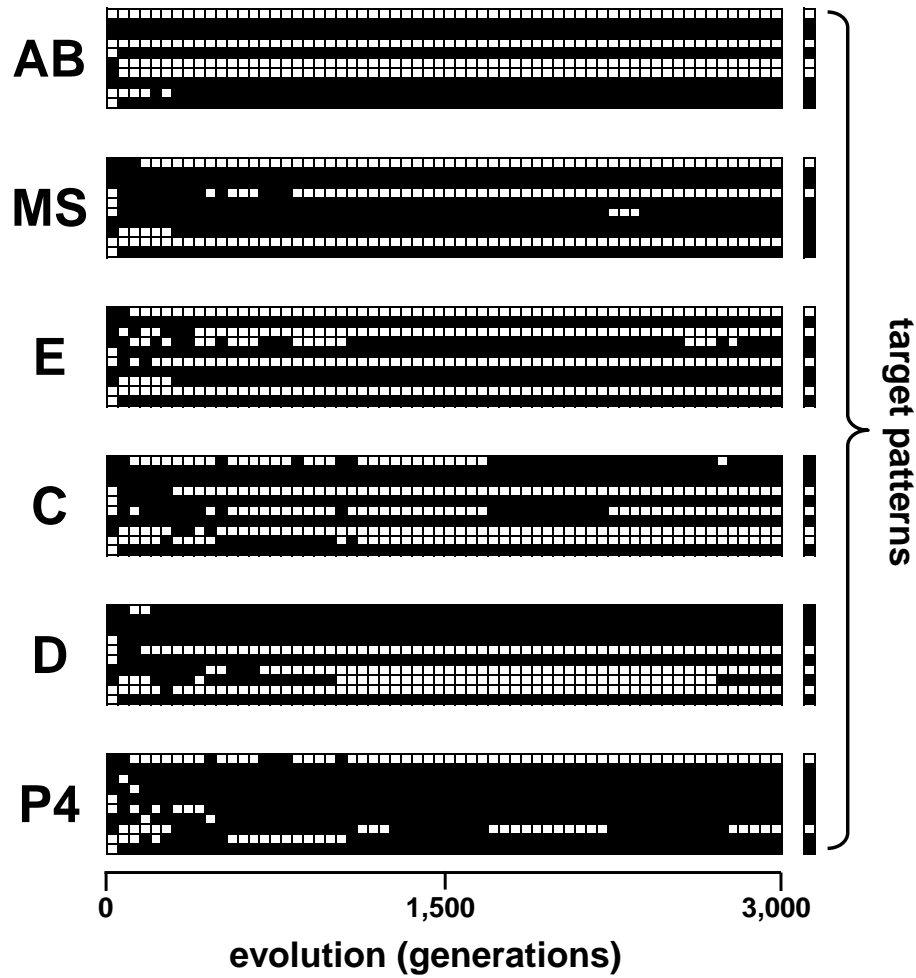


Fig. 7. Progress toward a solution in a typical evolutionary run. Snapshots of the system output are shown at intervals of 50 generations up to the 3,000th generation, at which point the system had achieved correct patterns of gene activation in all but one gene. Each of the large rows represents the evolution of gene activation of a single cell. Within a single cell's output, each individual row represents the activation of a particular gene. To increase clarity, genes with an activation greater or less than 0.5 are coloured black or white respectively.

References

1. B. Alberts, D. Bray, J. Lewis, M. Raff, K. Roberts, and J. D. Watson. *Molecular Biology of the Cell*. Garland Publishing Inc., New York, NY, 3rd edition, 1994.
2. P. J. Bentley, editor. *Creative Evolutionary Design*. Morgan Kaufman Publishers Inc., San Francisco, CA, 1999.
3. P. J. Bentley and S. Kumar. Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 25–43, Orlando, FL, 1999. Morgan Kaufmann.
4. J. Bongard and R. Pfeifer. Evolving complete agents using artificial ontogeny. In *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, pages 237–258. Springer-Verlag, Berlin, 2003.
5. S. Bornholdt. Modeling genetic networks and their evolution: A complex dynamical systems perspective. *Biological Chemistry*, 382:1289–1299, 2001.
6. H. de Jong. Modeling and simulation of genetic regulatory systems: A literature review. *Journal of Computational Biology*, 9(1):67–103, 2002.
7. F. Dellaert and R. D. Beer. A developmental model for the evolution of complete autonomous agents. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, editors, *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 393–401, Cambridge, MA, 1996. The MIT Press/Bradford Books.
8. R. Edwards and L. Glass. Combinatorial explosion in model gene networks. *Chaos*, 10(3):691–704, 2000.
9. P. Eggenberger. Evolving morphologies of simulated 3D organisms based on differential gene expression. In P. Husbands and I. Harvey, editors, *Fourth European Conference on Artificial Life*, pages 205–213, Cambridge, MA, 1997. The MIT Press/Bradford Books.
10. J. L. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
11. K. Fleischer. Investigations with a multicellular developmental model. In C. G. Langton and T. Shimohara, editors, *Artificial Life V: Proceedings of the Fifth International Conference on the Synthesis and Simulation of Living Systems*, pages 229–236, Cambridge, MA, 1996. The MIT Press/Bradford Books.
12. C. Furusawa and K. Kaneko. Origin of multicellular organisms as an inevitable consequence of dynamical systems. *The Anatomical Record*, 268:327–342, 2002.
13. N. Geard and J. Wiles. Structure and dynamics of a gene network model with RNA regulation. To be published in the Proceedings of the 2003 Congress on Evolutionary Computation., 2003.
14. F. Gruau. Automatic definition of modular neural networks. *Adaptive Behaviour*, 3(2):151–183, 1995.
15. P. Hogeweg. Shapes in the shadow: Evolutionary dynamics of morphogenesis. *Artificial Life*, 6:85–101, 2000.
16. F. Jacob and J. Monod. On the regulation of gene activity. *Cold Spring Harbor Symp. Quant. Biol.*, 26:193–211, 1961.
17. N. Jacobi. Harnessing morphogenesis. Cognitive Science Research Paper 423, School of Cognitive and Computing Sciences, University of Sussex, 1995.
18. S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford, UK, 1993.

19. H. Kitano. Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4(4):461–476, 1990.
20. G. Marnellos and E. Mjolsness. A gene network approach to modeling early neurogenesis in *Drosophila*. In R. B. Altman, A. K. Dunker, L. Hunter, and T. E. Klein, editors, *Pacific Symposium on Biocomputing '98*, pages 30–41, Singapore, 1998. World Scientific.
21. G. F. Miller. Artificial life as theoretical biology: How to do real science with computer simulation. Cognitive Science Research Paper 378, School of Cognitive and Computing Sciences, University of Sussex, 1995.
22. E. Mjolsness, D. H. Sharp, and J. Reinitz. A connectionist model of development. *Journal of Theoretical Biology*, 152:429–453, 1991.
23. S. Nolfi and D. Parisi. Growing neural networks. Technical Report PCIA-91-15, Institute of Psychology, CNR, Rome, 1991.
24. R. Serra, M. Villani, and A. Salvemini. Continuous genetic networks. *Parallel Computing*, 27:663–683, 2001.
25. M. Sipper. *Machine Nature: The Coming Age of Bio-Inspired Computing*. McGraw Hill, New York, NY, 2002.
26. R. Somogyi and C. A. Sniegoski. Modelling the complexity of genetic networks: Understanding multigenic and pleiotropic regulation. *Complexity*, 1:45–63, 1996.
27. K. O. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.
28. J. E. Sulston, E. Schierenberg, J. G. White, and J. N. Thompson. The embryonic cell lineage of the nematode *Caenorhabditis elegans*. *Developmental Biology*, 100:64–119, 1983.
29. The *C. elegans* Sequencing Consortium. Genome sequence of the nematode *C. elegans*: A platform for investigating biology. *Science*, 282:2012–2018, 1998.
30. B. Tonkes, J. Wiles, and J. S. Mattick. Controlling complexity in biological networks. Poster presented at the 11th International Conference on Intelligent Systems for Molecular Biology., 2003.
31. P. van Remortel, T. Lenaerts, and B. Manderick. Lineage and induction in the development of evolved genotypes for non-uniform 2D CAs. In *Proceedings of the 15th Australian Joint Conference on Artificial Intelligence*, volume 2557 of *Lecture Notes in Computer Science*, pages 321–332, London, UK, 2002. Springer-Verlag.
32. E. P. van Someren, L. F. A. Wessels, E. Backer, and M. J. T. Reinders. Genetic network modeling. *Pharmacogenomics*, 3(4):507–525, 2002.
33. J. Vohradský. Neural model of the genetic network. *The Journal of Biological Chemistry*, 276(39):36168–36173, 2001.
34. J. Vohradský. Neural network model of gene expression. *The FASEB Journal*, 15:846–854, 2001.
35. A. Wagner. Does evolutionary plasticity evolve? *Evolution*, 50(3):1008–1023, 1996.
36. L. Wolpert. *The Principles of Development*. Oxford University Press, Oxford, UK, 1998.
37. A. Wuensche. Genomic regulation modeled as a network with basins of attraction. In R. B. Altman, A. K. Dunker, L. Hunter, and T. E. Klein, editors, *Pacific Symposium on Biocomputing '98*, pages 89–102, Singapore, 1998. World Scientific.