Chapter 1

Evolving Gene Regulatory Networks for Cellular Morphogenesis

Timothy Rudge and Nicholas Geard

The ARC Centre for Complex Systems,
School of Information Technology and Electrical Engineering,
The University of Queensland, QLD 4072, Australia.
E-mail: {timrudge,nic}@itee.uq.edu.au

The generation of pattern and form in a developing organism results from a combination of interacting processes, guided by a programme encoded in its genome. The unfolding of this programme involves a complex interplay of gene regulation and intercellular signalling, as well as the mechanical processes of cell growth, division and movement.

In this study we present an integrated modeling framework for simulating multicellular morphogenesis that includes plausible models of both genetic and cellular processes, using leaf morphogenesis as an example. We present results of an experiment designed to investigate the contribution that genetic control of cell growth and division makes to the performance of a developing system.

1.1 Introduction

The generation of pattern and form in a developing organism results from a combination of interacting processes, guided by a programme encoded in its genome. The unfolding of this programme involves a complex interplay of gene regulation and intercellular signalling, as well as the morphogenetic processes of cell growth, division and movement [20]. Recently, computers have enabled these multi-scale developmental systems to be simulated, revealing new insights into the emergence of pattern and form [9; 15].

In this study we present an integrated modeling framework for simulating multicellular morphogenesis that includes plausible models of both genetic and cellular processes, using leaf morphogenesis as an example. Leaf forms display a wide variety of morphological features, the development of which provide excellent examples of robust control of shape formation. We focus here on the role played by genetic control of cell growth and division orientation in the generation of specific shapes, and how this interacts with the physical constraints on cell shape.

The format of this report is as follows: Section 1.2 presents background material on leaf morphogenesis and previous computational models of gene regulation and development. Section 1.3 describes our simulation framework, which consists of a gene network model, a physical cell model, a procedure for coupling these two components, and an evolutionary search algorithm for investigating model parameters. Section 1.4 presents initial results obtained using this simulation framework and Section 1.5 concludes with a discussion of future directions for this research.

1.2 Background

1.2.1 *Leaf Morphogenesis*

Morphogenesis, the formation of shapes and structures in plants and animals, occurs by three processes: 1. Tissue growth; 2. Cell movement; 3. Cell death (apoptosis) [2]. Active cell movement does not occur in plants and so morphogenesis is coordinated by tissue growth – determined by cell shape, growth, and proliferation – and cell death. Variation in these behaviours across tissues and over developmental time causes the development of specific forms. The term *patterning* is applied to the coordinated differential expression of genes over space and time. It is these gene expression patterns that give rise to the variation in cell behaviour that drives morphogenesis.

Patterning provides positional information that guides cell behaviour and although cell lineage also plays some role, it seems this positional information is of primary importance in plant development [3; 7]. The relationship between patterns of gene expression and the specification of tissue and organ shape is not well characterised. [2] cite the difficulty in measuring morphogenetic effects and the need for quantitative analysis as possible reasons for this gap. Another difficulty is understanding the tight coupling of morphogenesis and patterning: the patterns develop along with, and are embedded in the forms to which they give rise.

Formation of leaf shape is tightly regulated, and evidence exists for

both cell-division dependent and cell-division independent regulatory mechanisms [7]. That is, final leaf form and size is to some extent independent of both cell number and cell size. There is also evidence that overall leaf shape is unaffected by cell division orientation [7]. This suggests regulation of cell behaviour depends on feedback of organ-level information [19]. The nature of this regulation is at present unknown. However, hormone (e.g. auxin) transport has been implicated in organ shape regulation [7; 12], and [11] suggest that hormones may also regulate organ shape by affecting cell expansion and/or by modulating the cell cycle. Noting that leaf size is dependent on whole-plant physiology, [19] suggests that a source-sink relationship within the plant (e.g. of nutrients) might limit leaf size.

[21] and [11] present results of experiments in which cell division rates, cell division orientation, and cell growth rates were perturbed both locally and across the leaf. These results provide evidence for the involvement of the above mentioned processes in regulating leaf morphogenesis, but how these multiple mechanisms interact and their relative importance are still unknown.

1.2.2 *Previous Models*

[13] identifies three categories of plant development models, focusing on plant architecture, individual organs, and the underlying mechanics of gene regulation, respectively. The first of these is well established, with L-systems being the dominant modelling framework; the latter two areas are still active areas of research – [13] and [2] provide an overview of recent developments. These categories occupy very different temporal and spatial scale ranges, and a full understanding of development requires the integration of multiple scales.

Developmental issues have been addressed by the Artificial Life research community [18]. An early attempt to integrate multiple scales of developmental mechanism into a single model included cells with complex internal dynamics that communicated with each other via chemical and electrical signals as well as physical interactions [6]. One of the findings of this study was that, while the multiple mechanisms enabled the robust production of interesting phenotypes, it also made the design of *specific* phenotypes more difficult. Later research demonstrated that this difficulty could be addressed by using a representation of the regulatory network that could be artificially evolved [4].

[9] also used a model that combined mechanisms at multiple scales of description – gene regulation, development and evolution – to investigate the interactions between evolutionary dynamics and morphology. Her evolutionary process was aimed at maximising cell type diversity, rather

than achieving a specific morphological shape. However, she found that some features of morphology, such as engulfing, budding and elongation, were relatively 'generic', that is, they appeared 'for free' in systems that satisfied certain prerequisites.

The formation of patterns across a fixed field of autonomous cells has also been studied in some detail [16; 17]. Here too, it was found that patterns were a common emergent feature of interacting gene networks, although again, selection was for 'pattern complexity' rather than a specific phenotypic target. Later research [15] focused on the issue of how pattern formation processes interact with growth processes, specifically with reference to the evolution of tooth development.

In general, the models mentioned above do not include individual cell morphology: Cells are represented as circles or squares, of equal or varying size. The model used by [9] does support anisotropic cell shapes; however, these result from the algorithm used to calculate cell boundaries, rather than reflecting anisotropy in the underlying growth process. In plants, cell behaviour is frequently anisotropic, with the axes of both growth and division under a degree of cellular and genetic control. As described above, control of cell morphology is intimately connected with the production of leaf form, therefore a detailed and flexible model of cell shape is of fundamental importance in any approach to modeling leaf morphogenesis.

1.3 The Simulation Framework

Our integrated model of plant morphogenesis brings together plausible representations of cell shape, genetic regulation, and cell-cell signalling. Cell shape is determined by growth and division activity as well as external physical forces, and the combination of the shapes of all cells determines the overall phenotypic form. Each cell is autonomous and its behaviour is regulated by its own copy of the organism's gene network, which also responds to signals received from neighbouring cells. The genetic network thus indirectly specifies phenotypic morphology.

As noted earlier, due to the complexity of cellular developmental systems, a search and optimisation approach is favoured when examining their properties *in silico*. We have chosen to use an evolutionary algorithm to search for systems with particular shape formation capabilities. Our approach is to decide on a target phenotypic shape, specify initial conditions, and then artificially evolve gene networks which come close to producing the desired shape.

In the following we describe the primary components of the model: the genetic component, consisting of a network embedded within each cell, and

the spatial model, consisting of an arrangement of cells that constitutes the phenotype. Following that, the coupling between these two components is described. Finally, the evolutionary algorithm used to explore the parameter space of these networks is outlined.

1.3.1 *The Genetic Component*

In this study, we used a dynamic recurrent gene network (DRGN) model for the genetic component of the framework [8]. The DRGN model is based on a widely studied class of artificial neural network models known as recurrent neural networks [5], and has previously been used to investigate the generation of developmental cell lineages [8]. An advantage of a recurrent network representation is that it enables the model to express a complex range of gene interactions while abstracting away from the specific biological processes that underly those interactions.

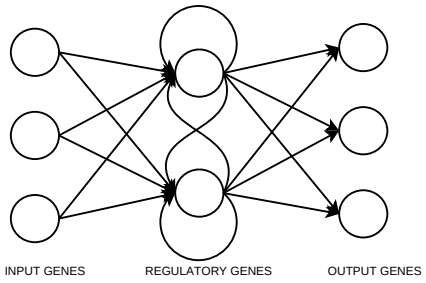

Fig. 1.1 The structure of the DRGN model. The network is partitioned into three categories: input genes that detect the presence of morphogens produced by other cells or the environment; regulatory genes that interact with one another to perform the computational tasks of the cell; and output genes producing morphogens that can be transmitted to other cells or that trigger events such as growth and division.

In the DRGN model, a genetic system is defined as a network of N interacting nodes (see Figure 1.1). Depending on the level of abstraction, each node can be considered to represent either a single gene, or a cluster of co-regulated genes. In this study we generally consider a node to be a equivalent to a single gene. The activation state of each node is a continuous variable in the range $[0, 1]$, where 0 represented a completely inactive gene and 1 a fully expressed gene. Nodes can be divided into three classes: input genes that detect the presence of morphogens; regulatory genes that interact with each other to carry out the computational task of the network; and output genes that produce morphogen signals.

The network is updated synchronously in discrete time steps. To capture the potential complexity of the interacting factors involved in gene expression, we have used a network in which each input gene is connected to each regulatory gene, all regulatory genes are connected to each other and themselves, and each regulatory gene is connected to each output gene. Thus an individual link in the network does not necessarily represent a direct physical interaction, but rather the degree of influence that the expression of the source gene at time t has on the expression of the target gene at time $t+1$.

These interactions can be summarised in a weight matrix, in which the entry at row i, column j specifies the influence that gene j has on gene i. These entries may be positive or negative, depending on whether the product of gene j is an activator or a repressor in the regulatory context of gene i. A zero entry indicates that there is no interaction between the two genes. The inclusion of self-connections (i.e. from node i to node i) allows for the possibility of genes influencing their own regulation.

The state of the network is updated synchronously, with the activation of node i at time $t+1$, $a_i(t+1)$, given by

$$a_i(t+1) = \sigma\big(\sum_{j=1}^{N_r} w_{ij} a_j(t) - \theta_i\big) \tag{1.1}$$

where N_r is the number of regulatory nodes, w_{ij} is the level of the interaction from node j to node i, θ_i is the activation threshold of node i, and $\sigma(.)$ is the sigmoid function, given by

$$\sigma(x) = \frac{1}{1+e^{-x}} \tag{1.2}$$

1.3.2 *The Cellular Component*

We use a 2-dimensional spatial model of the cellular arrangement. This is based on linear cell boundary elements (walls), which are modelled as elastic springs. The approach is similar to that of [10], however we also consider some more complex cell dynamics such as anisotropic growth. Cell-cell signalling is considered in the form of chemical diffusion, as in [6]. This approach has previously been used to examine rule-based control of plant morphogenesis [14].

Cell: The genome of our artificial organism is represented as a DRGN. Each cell is defined by its DRGN, a set of dynamic state parameters and a closed boundary. The DRGNs contained by each cell in a phenotype have identical structure and weights, reflecting the genetic homogeneity of an individual organism. The activation levels of the DRGN nodes in each cell,

however, are independent and represent the variation in gene expression across the phenotype.

The cell state parameters include passively received information such as morphogen levels and cell volume, and behavioural states like growth rate and morphogen production rates. As part of their state, the cells also maintain polarity vectors that are used to direct anisotropic growth, to orient the division plane, and to asymmetrically divide the cells morphogens between its daughters, according to the behavioural state parameters.

The state of the cell determines its behaviour at any point in time. Cell dynamics are expressed as the transformation of cell state parameters to proceed to a new state. Behavioural states are transformed by the DRGN, with the inputs and outputs of the DRGN defined by a fixed mapping onto the cell state parameters. The passive state parameters are transformed by physical simulation of the cells' environment, including its own boundary shape and interactions with neighbours.

Spatio-Mechanical Model: The boundary of the cell describes its shape, and is decomposed into a set of walls. Each wall is the interface between two cells. Morphogens diffuse from one cell to the other via the wall, providing a cell-cell signalling mechanism. The walls are considered to be two linearly elastic elements (springs), one for each adjacent cell, bound together at the end points (vertices). Each of the adjacent cells influences the properties of only one of these springs. Each spring has stiffness K and natural length L_n determined from the state parameters of the appropriate cell [14].

Each cell exerts a turgor force perpendicular to each of its walls in an outward direction with respect to the cell, extending the springs, which then exert an opposing tension force. At each time-step these simulated forces are accumulated at the vertices, and the vertex positions are adjusted to find the equilibrium configuration.

Cell growth is achieved by increasing the natural lengths of each cells' springs to varying degrees (see [14] for details). Division consists of inserting a dividing wall across the centre of the cell, and redefining the daughter cell boundaries. When a cell divides, its DRGN (including current node activation levels) is copied into the two daughter cells.

1.3.3 *Genotype-Phenotype Coupling*

The system integrates multiple scales of model into a single framework. Figure 1.2 shows an overview of the way in which the levels of the model interact. Starting at the micro level, the DRGN transforms the cell state. The cell state is expressed as local behaviours such as growth, which then affect the entire phenotype via simulated mechanical forces and diffusion

processes. This global effect is then transduced back into local information to each cell, and from there transformed into micro level input to the DRGN.

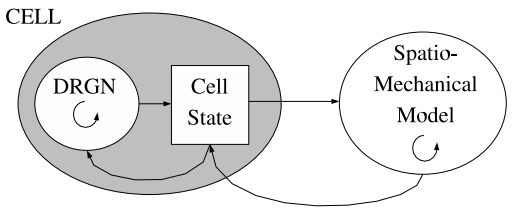

Fig. 1.2 Scheme of interactions between different levels in the model, from microscopic (left) to macroscopic (right). Circular arrows indicate faster time scale processes running multiple time steps between cell state updates.

The flow of control is therefore from micro to macro level, and the flow of information or feedback is from macro to micro. The nature of the coupling between information feedback and phenotypic output is ultimately determined by the structure and weights of the DRGN, as well as the mapping from DRGN to cell state – i.e. the genotype.

The dynamics of cell behaviour, such as growth and division, gene expression, and transmission of mechanical forces, occur on very different time scales. In general, variation in cell behaviour occurs most slowly and equilibration of forces occurs most quickly. We assume that mechanical equilibrium is reached instantaneously when relevant parameters such as growth rate change.

The DRGN can be used to model genetic regulation on several levels. Each node may represent a single gene or a cluster of genes, and each node update may represent one or many regulatory events. In order to incorporate this flexibility we allow the DRGN to update multiple times before affecting the cell state. The procedure that produces a cellular phenotype from the DRGN genotype is thus:

(1) Determine cell states from initial conditions
(2) Map DRGN inputs from cell states
(3) Update DRGN by some number of time steps
(4) Map cell states from DRGN outputs
(5) Compute cell shapes and morphogen diffusion
(6) Repeat from 2 until stopping condition met

The stopping condition may be chosen arbitrarily according to the ex-

periment. We used a maximum number of time steps of 250 in our experiments.

1.3.4 *The Evolutionary Component*

The evolutionary component, which enables a population of DRGNs to be artificially "evolved" towards some particular target, serves two purposes. At a methodological level, it provides a useful machine learning technique for searching the parameter space of networks. At a theoretical level, it facilitates questions about the evolutionary dynamics of morphogenesis [9].

A simple evolutionary search strategy called the 1+1 ES was used [1]. Initially, a single DRGN was generated with weights randomly drawn from a Gaussian distribution with mean 0 and standard deviation 4. This DRGN was used to develop a phenotype, as described in Section 1.3.3. A fitness value for this phenotype was calculated as described in Section 1.4.1 below and stored. A new DRGN was derived from the existing DRGN by adding Gaussian noise (mean 0, standard deviation 0.01) to each of the node interactions. A new phenotype was developed and evaluated and the fitness value for the modified DRGN was compared to that of the original DRGN. The DRGN producing the phenotype with the greatest fitness was retained and used as the basis for the creation of a further new DRGN. This process was repeated until the stopping conditions were met. We used a maximum number of generations of $15,000$ in our experiments.

1.4 Initial Experiments

To investigate the role of genetic control of growth and development in morphogenesis, we ran three sets of comparative evolutionary trials: random growth and division orientation, regular growth and division orientation and genetically controlled growth and division orientation. We set the DRGN the task of generating a circular shaped final phenotypic form.

(1) **Random orientation** In the first set of trials, there was no control of growth and division orientation – they were each chosen randomly at each time step. Only one output node was utilised, the morphogen controlling the decision to grow and divide.
(2) **Regular orientation** In the second set of trials, the orientation of growth and division of each cell was chosen to be opposite to that of its parent cell – giving alternating axial and lateral growth and division each generation. The DRGN was not able to change this sequence of orientations; however, it was able to make use of the regularity of the

predetermined sequence.

(3) **Controlled orientation** In the final set of trials, two additional output nodes were used, producing morphogens that controlled the orientation of growth and division respectively. Therefore the DRGNs had the capability to coordinate the two processes.

1.4.1 *Method*

DRGN Coupling: Three inputs were provided to the network. The first input responded to the concentration of a morphogen that was initialised to a concentration of 1.0 in the initial cell, and was not produced after that. Therefore, as the volume of the phenotype increased due to growth and division, the concentration of this morphogen decreased. The second and third inputs responded to morphogens related to the position of the cell. These were externally supplied as the (x, y) position of the cell centre. This may be considered as incorporating information supplied by underlying cell layers.

The phenotypes were initialised as a single unit square cell with unit morphogen concentration, and DRGN outputs p_j mapped to cell behaviour as follows:

- If $p_0 > 0.5$ then set growth rate to 0.2, and divide if volume > 2.
- If $p_0 \leq 0.5$ then set growth rate to 0 and do not divide.
- If $p_1 > 0.5$ set division orientation to *axial* otherwise set to *lateral*.
- If $p_2 > 0.5$ set growth orientation to *axial* otherwise set to *lateral*.

Fitness function: The task used for this study was to evolve a DRGN capable of generating a circular arrangement of cells of a given radius. Fitness was calculated for each phenotype at each time step based on the current cell arrangement – specifically, the absolute distance of each exterior (marginal) cell from the centre of mass of the phenotype, given by:

$$r_i = |\mathbf{x_i} - \mathbf{c}| \tag{1.3}$$

for cell i, where $\mathbf{x_i}$ is the cell's centre of mass, and $\mathbf{c}$ is the centre of mass of the whole phenotype. The error of the phenotype from a circle radius R at any time point, is calculated from the distance of each cell from the circle $dr_i = |r_i - R|$:

$$\epsilon(t) = \bar{dr}(t) + \sqrt{\frac{1}{N(t)} \sum_{i=1}^{N(t)} \left(dr_i(t) - \bar{dr}(t)\right)^2} \tag{1.4}$$

where the sum is over the N exterior cells. The first term is the mean distance error, and the second is the standard deviation in the distance error. We used a target radius of 5 units in our experiments. The overall phenotypic fitness was calculated from the cumulative error over all time steps $\{1, 2, 3...T\}$ and scaled by a constant factor such that the maximum possible fitness is approximately 1.0:

$$f = \frac{625}{1 + \sum_{t=1}^{T} \epsilon(t)} \tag{1.5}$$

Summary: In summary, three sets of evolutionary trials were run, each corresponding to one of the control conditions described in Section 1.4. Each evolutionary trial was run for up to $15,000$ generations, with snapshots of the best phenotype being recorded at 500 generation intervals. In each generation, the DRGN was run for 250 developmental time steps, with it's fitness evaluated over this period as described above.

1.4.2 *Results*

The DRGNs that had explicit control of the growth and division orientation were able to generate considerably more accurate phenotypes than the those supplied with either a regular or random sequence of orientations (Table 1.1).

Table 1.1

Orientation	**Maximum Fitness**
Random	0.314
Regular	0.555
Controlled	0.642

The evolutionary history of the most successful evolutionary trial from the Controlled set displays a level of continual innovation typical of highly evolvable systems (Figure 1.3). By contrast, the most successful trials from the Regular and Random sets (not shown), reached their peak fitness early (around generation 4,000), and failed to improve any further. The phenotypes produced at different stages of evolution provide some clues to explain these differences (Figure 1.4).

(1) **Random orientation (Figures 1.4(a)- 1.4(d)):** Regulation of size appears relatively early, and is consistent throughout the course of evolution. However control of phenotypic shape has not evolved. The randomness of the sequence of orientations prevents the DRGN from

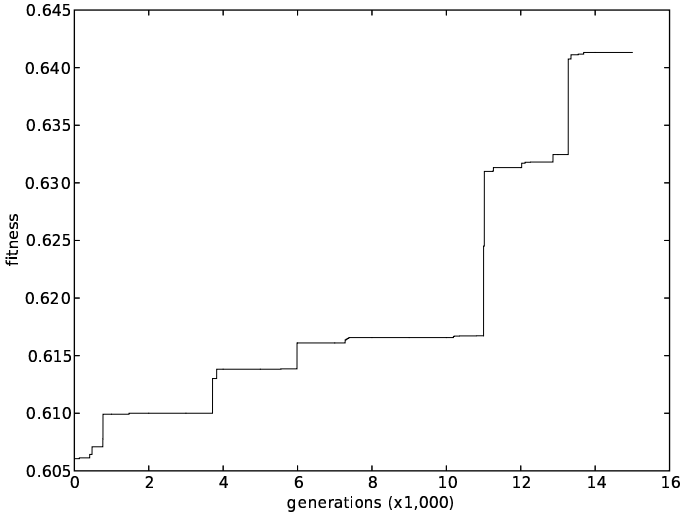

Fig. 1.3 The evolutionary history of the fittest system found in the Controlled set of trials, in which the DRGN had explicit control of the orientation of growth and division.

being able to successfully coordinate the development of a stable shape. The examples shown here represent only one instance of a set of possible outcomes for a given DRGN.

(2) **Regular orientation (Figures 1.4(e)- 1.4(h)):** When the DRGN is able to rely on a regular sequence of growth and division orientations, greater control of phenotypic shape is achieved. Very early in this evolutionary trial, a strategy emerged in which a group of growing cells is surrounded by non-growing cells. However, while ensuring a reasonably circular phenotype, this approach proves too strong a constraint, limiting any further improvement.

(3) **Controlled orientation (Figures 1.4(i)- 1.4(l)):** With full control over division and growth orientations, the evolutionary algorithm was able to explore a much broader range of developmental possibilities. In the example shown, the DRGNs found early in the evolutionary history developed by first growing and dividing laterally and then switching to axial division, resulting in the phenotype fanning out. The largest jump in fitness (Figure 1.3, around generation 11,000) occurred when a DRGN was discovered in which the fanning out process was inhibited by a cap of quiescent cells. The resulting "stem and bud" arrangement

was refined in successive stages of evolution by increasing the roundness of the bud, and reducing the length of the stem.

1.5 Discussion and Future Directions

In all three sets of trials, DRGN evolved that were capable of controlling phenotype size. With full DRGN control over development a significant degree of shape control evolved, using a variety of developmental approaches. With regular cell growth and division, DRGNs were able to control the phenotype shape to a similar extent, but the range of strategies for doing this was more limited. With random growth and division, little control of phenotypic shape emerged.

While preliminary, these results suggest that positional information only provides sufficient information to enable generation of stable phenotypic forms in the presence of predictable growth and division orientation. It would appear that the claim by [7], that leaf shape is, to a degree, independent of division orientation, requires the presence of more complex cell-cell signalling than simple positional cues. One strong possibility is that cell-cell communication plays a vital role in the robust development of form. Future work will involve the investigation of inductive interactions between cells and how this additional level of communication may facilitate more robust development.

Acknowledgments

This study was funded by the ARC Centre for Complex Systems (`http://www.accs.edu.au`). We thank Jim Hanan and Janet Wiles for stimulating discussions, helpful suggestions and guidance in carrying out this research.

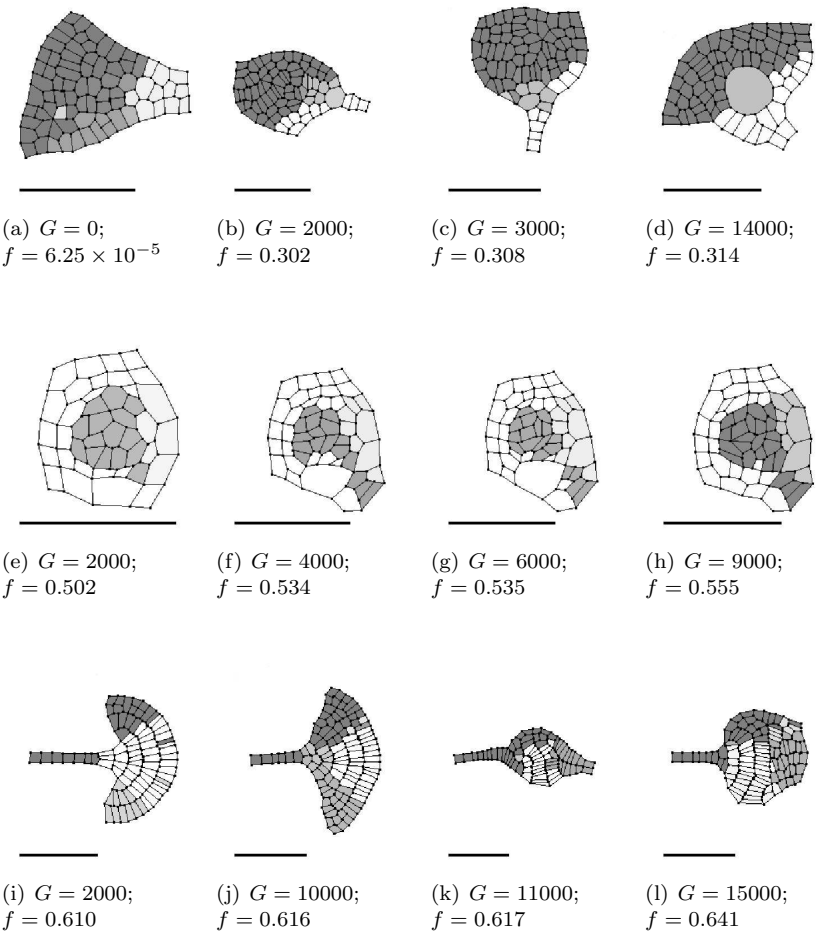

(a) $G = 0$; $f = 6.25 \times 10^{-5}$ (b) $G = 2000$; $f = 0.302$ (c) $G = 3000$; $f = 0.308$ (d) $G = 14000$; $f = 0.314$

(e) $G = 2000$; $f = 0.502$ (f) $G = 4000$; $f = 0.534$ (g) $G = 6000$; $f = 0.535$ (h) $G = 9000$; $f = 0.555$

(i) $G = 2000$; $f = 0.610$ (j) $G = 10000$; $f = 0.616$ (k) $G = 11000$; $f = 0.617$ (l) $G = 15000$; $f = 0.641$

Fig. 1.4 Fittest phenotypes at key stages (generation G) in artificial evolution, where f is fitness: [1.4(a)-1.4(d)] Random growth/division, [1.4(e)-1.4(h)] Regular growth/division, [1.4(i)-1.4(l)] DRGN control. Shading shows DRGN output on a grey scale, white(0) to dark grey(1): [1.4(a)-1.4(h)] cell growth and division trigger, [1.4(i)-1.4(l)] division orientation. Scale bar is 10 units.

Bibliography

T. Bäck, D. Fogel, and Z. Michalewicz. *Handbook of Evolutionary Computation*. Oxford University Press, Oxford, UK, 1997.

E. Coen, A.-G. Rolland-Lagan, M. Matthews, J. A. Bangham, and P. Prusinkiewicz. The genetics of geometry. *Proceedings of the National Academy of Science, USA*, 101:4728–4735, 2004.

L. Dolan and K. Roberts. Plant development: pulled up by the roots. *Current Opinion in Genetics and Development*, 5:432–438, 1995.

P. Eggenberger. Evolving morpohologies of simulated 3D organisms based on diffferential 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.

J. L. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.

K. Fleischer and A. H. Barr. A simulation testbed for the study of multicellular development: The multiple mechanisms of morphogenesis. In C. Langton, editor, *Artificial Life III*, pages 389–416, Redwood City, CA, 1994. Addison-Wesley.

A. J. Fleming. The mechanism of leaf morphogenesis. *Planta*, 216:17–22, 2002.

N. Geard and J. Wiles. A gene network model for developing cell lineages. *Artificial Life*, 11(3):249–268, 2005.

P. Hogeweg. Shapes in the shadow: Evolutionary dynamics of morphogenesis. *Artificial Life*, 6:85–101, 2000.

H. Honda, H. Yamanaka, and M. Dan-Sohkawa. A computer simulation of geometrical configurations during cell division. *Journal of Theoretical Biology*, 106:423–435, 1984.

S. Kessler and N. Sinha. Shaping up: the genetic control of leaf shape. *Current Opinion in Plant Biology*, 7:65–72, 2004.

G. T. Kim, H. Tsukaya, Y. Aito, and H. Uchiyama. Change in the shape of leaves and flowers upon overexpression of cytochrome p450 in arabidopsis. *Proc Natl Acad Sci USA*, 96:943309437, 1999.

P. Prusinkiewicz. Modeling plant growth and development. *Current Opinion in Plant Biology*, 7:79–83, 2004.

T. Rudge and J. Haseloff. A computational model of cellular morphogenesis

in plants. To appear in *VIIIth European Conference on Artificial Life*, Canterbury, UK, 2005.

I. Salazar-Ciudad and J. Jernvall. How different types of pattern formation mechanisms affect the evolution of form and development. *Evolution and Development*, 6:6–16, 2004.

I. Salazar-Ciudad, S. A. Newman, and R. V. Solé. Phenotypic and dynamical transitions in model gene networks I. Emergence of patterns and genotype-phenotype relationships. *Evolution and Development*, 3:84–94, 2001.

R. V. Solé, I. Salazar-Ciudad, and J. Garcia-Fernández. Common pattern formation, modularity and phase transitions in a gene network model of morphogenesis. *Physica A*, 305:640–654, 2002.

K. O. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.

H. Tsukaya. Organ shapen and size: a lesson from studies of leaf morphogenesis. *Current Opinion in Plant Biology*, 6:57–62, 2003.

L. Wolpert. *The Principles of Development*. Oxford University Press, Oxford, UK, 1998.

J. Wyrzykowska, S. Pien, W. H. Shen, and A. J. Fleming. Manipulation of leaf shape by modulation of cell division. *Development*, 129:957–964, 2002.