DIVISION OF COMPUTER SCIENCE

Artificial Evolution: Modelling the Development of the Retina

A G Rust
R Adams
S George
H Bolouri

Technical Report No.255

September 1996

Artificial Evolution: Modelling the Development of the Retina

Alistair G Rust
Department of Computer Science &
Engineering Research and Development Centre

Rod Adams
Department of Computer Science

Stella George
Department of Computer Science

Hamid Bolouri
Engineering Research and Development Centre

University of Hertfordshire, AL10 9AB, UK
Tel : +44 1707 284661 Fax : +44 1707 284185
E-mail : a.g.rust, r.g.adams, s.j.george, h.bolouri@herts.ac.uk
WWW: http://erdc.herts.ac.uk/NAGweb/NAGhome.html

Computer Science Technical Report No.255
ERDC Technical Report ERDC/1996/0015

September 13, 1996

Abstract

The evolution of neural systems relies on the repeated modification of developmental programmes contained within genes. This paper proposes that to efficiently investigate artificial evolution, developmental processes must first be characterised and encoded. The development processes themselves are inherently simple, representing a form of emergent self-organisation. A model of artificial neural structure self-organisation is presented, where development occurs across different levels of scale and development processes are divided into phases. Within this framework, a Developmental Artificial Neural Network (DANN) has been implemented to model the growth of neuron-neuron connections in both 2D and 3D. The development of a biologically inspired, artificial retina is being investigated. This is believed to be the first reported 3D DANN.

1 Introduction

1.1 Evolution and Development

It is generally recognised that the successful evolution of non-trivial neural systems requires non-linear genotype to phenotype mappings [Kitano, 1990, Nolfi and Parisi, 1991]. Current evolutionary neural network models however, use ad-hoc, non-linear mappings [Boers et al., 1993, deGaris, 1995, Gruau, 1994] and have little in common with their biological counterparts.

This paper proposes an alternative, bottom-up approach. Mechanisms of biological evolutionary adaptation (gene duplications, followed by mutations of the duplicates) are intimately linked with the mechanisms of gene expression which lead to the development of neural systems. Evolutionary changes are due to the genetic modifications of the developmental processes. Thus, artificial evolution of neural systems requires developmental processes to be appropriately genetically encoded.

Biological developmental rules can be modelled using a variety of algorithmic techniques, for example, partial differential equations [Fleischer, 1995], L-systems [Boers et al., 1993, Vaario, 1994] and cellular automata [deGaris, 1995]. Our aim is to find the least complex, computational rule sets which allow growth of large 3D networks that mimic biological structure and function. Simulated networks can then be compared against actual biological networks. Minimal rule sets also provide jumping-off points to efficiently investigate artificial evolution.

The development of the mammalian retina is being modelled in this work, since the structure and function of the retina are biologically well documented, and does not appear to require learning. This enables the developed artificial neural systems development models to be directly compared against biological systems and simplifies the design of fitness functions for the evolutionary exploration of developmental parameters [Rust and Bolouri, 1996]. If it is possible to mimic the development of a retina, this may enable the development techniques to be used to create other neural structures. The long term objective is to incrementally develop a shape recognition system inspired by the primate biological visual system [Rust and Bolouri, 1996].

1.2 Self-Organisation

It is believed that the creation of the precise and intricate form of the brain is governed by simple developmental rules [Stryker, 1994, von der Malsburg, 1995]. These rules encompass the genetic encoding of structure and the modification of structure by neural activity. The interactions of these simple rules represents developmental self-organisation [von der Malsburg, 1995].

Self-organisation of structure within biological systems occurs at several levels of scale. One such division of structure self-organisation is illustrated in Figure 1. Here there are 3 levels: self-organisation of neuron structure, self-organisation of neuron connectivity, and self-organisation of neural modules. The self-organising processes are equally complex across all levels of scale. It is simply the level of detail (molecules, neurons, modules) which changes.

This document is only concerned with the self-organisation of structure across the last 2 levels illustrated in Figure 1. It is postulated that the same rules of self-organisation are applicable to both the development of neuron connectivity and the grouping of neural modules, suggesting that there is self-similarity across scales. Development is modelled as a 3-stage process, where each stage is a collection of abstract biological processes.

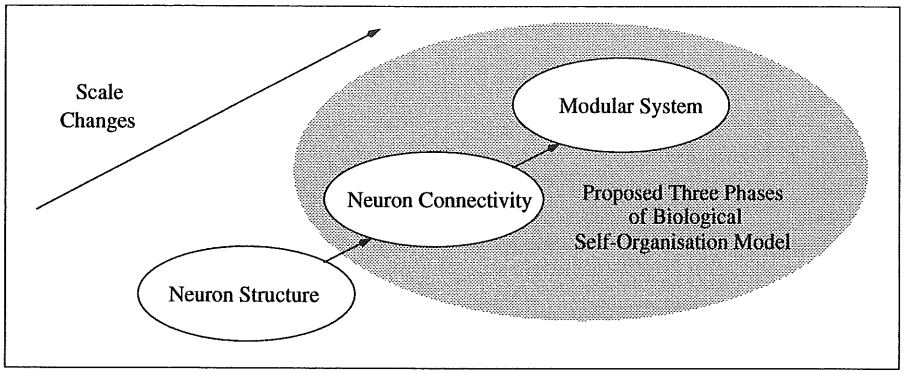

Figure 1: Self Organisation Across Scales

The document is organised as follows. The second section presents brief summaries of neural development and biological retinas. Developmental artificial neural networks are summarised in Section 3. Section 4 introduces *the three phases of biological self-organisation* model along with starting assumptions. The modelling environment, growth rules and a description of the implementation are discussed in Sections 5 and 6. Section 7 presents the results of modelling an artificial retina, which are discussed in Section 8 before conclusions are drawn.

2 Biological Reviews

The following section presents overviews of neural development and biological retinas. More detailed reviews may be found in [Rust and Bolouri, 1996].

2.1 A Brief Review of Neural Development

During neural development [Stryker, 1994]:

> *Neurons are generated in appropriate numbers at appropriate times; they migrate to appropriate positions; they send out axons that find appropriate paths to their target structures; and the axon terminals recognise the correct cell types and even the right general region in the target structures.*

These processes occur without electrical activity in neurons and appear to be driven by chemical events [Goodman and Shatz, 1993,Stryker, 1994]. It is believed that genes alone cannot provide the 'blueprint' for the final structure as this process can only create highly regular architectures [von der Malsburg, 1995] and that the chemical environment plays a critical role in the development process.

Once a neuron has reached its final position, it sends out axons and dendrites, collectively termed neurites [Alberts et al., 1994]. A neurite extends with a growth cone on the end of its fibre, which is covered in microscopic hairs called filopodia [Alberts et al., 1994]. Filopodia are extremely chemically sensitive and guide the growth cone to specific chemical targets by probing the immediate environment.

As a growth cone nears the point of termination it may be influenced by chemical gradients of neurotrophic factors produced from the target neurons [Alberts et al., 1994]. Not only may these factors provide the chemical recognition [Alberts et al., 1994] but they also control the survival of the subsequent axon connection and ultimately the neurons themselves [Purves and Lichtman, 1985]. Self-organisation is an emergent property of the simple, local interactions between neurites and chemical gradients in their environment.

The processes of genetic growth encoding and chemical gradients are able to create a skeletal framework but it is patterns of neural activity which subsequently produce the final, complex and optimised circuitry [Alberts et al., 1994,Shatz, 1994].

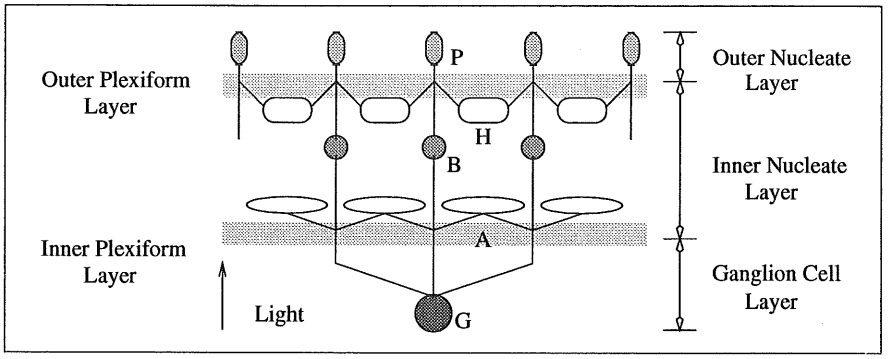

Figure 2: A simplified diagram of the retinal architecture. Key : P-photoreceptor, B-bipolar, H-horizontal, A-amacrine and G-ganglion Cells

2.2 An Overview of the Retina

The retina is one of the most studied organs in vertebrates [Dowling, 1992] and is described as having a 'stereotyped architecture' [Van Essen and Anderson, 1995] which is 'highly ordered anatomically' [Dowling, 1992].

Five major cell types are found in the retina - photoreceptor, bipolar, horizontal, amacrine and ganglion cells. Each cell type may be further divided into subclasses, for example there are two types of photoreceptors; rods and cones. These cells are arranged into three cellular layers - the Outer Nucleate Layer (ONL), the Inner Nucleate Layer (INL)and the Ganglion Cell Layer (GCL). Synaptic connections between cells occur in two distinct plexiform layers (Inner and Outer). The architecture of the retina is shown in Figure 2. In terms of information processing, experimental evidence has shown that the OPL processes static, spatial elements of images whilst the IPL deals with transient or temporal aspects [Dowling, 1992].

The topology of the retina is governed by simple interconnection rules :

1. There is a distinct vertical arrangement : Photoreceptors → Bipolar → Ganglion cells. Therefore, receptors cannot directly contact ganglion cells, with the bipolar cells acting as inter-layer links.
2. Lateral interactions are transmitted by horizontal and amacrine cells but within the respective plexiform layers.
3. At each plexiform layer, the synapses from three cell types interact; photoreceptors, bipolar and horizontal cells in the outer layer, and bipolar, ganglion and amacrine cells in the inner layer.

3 Developmental Artificial Neural Networks (DANNs)

3.1 The Current State of DANNs

There has been recent interest in incorporating biological development in the design and construction of artificial neural networks. There is however a great degree of variation in the detail with which biological processes are modelled in such DANNs [Rust and Bolouri, 1996]. There is a trade-off between the degree of biological plausibility and the success with which complex networks are evolved. Key attributes and findings of DANNs are:

1. Developmental rules can successfully model individual neural growth but an artificial environment is required to allow cell-cell and cell-substrate interactions [Dellaert and Beer, 1994, Fleischer, 1995, Jakobi, 1995, Nolfi and Parisi, 1991, Vaario, 1994]. Further, simply mapping directly from a cell genotype to a phenotype does not capture the complex, non-linear interactions that occur during neural development [Kitano, 1990, Nolfi and Parisi, 1991].
2. Cells or neurons can be modelled as self-contained elements which may be influenced both by internal and/or external factors [Fleischer, 1995, Jakobi, 1995, Vaario, 1994].

3. Modularity, symmetry breaking and scalability can be incorporated in the development process [Boers et al., 1993, Dellaert and Beer, 1994, Fleischer, 1995, Gruau, 1994, Kitano, 1990].

4. In a number of DANNs no learning occurs as the network's function is encoded within the development process [deGaris, 1995, Gruau, 1994, Jakobi, 1995, Nolfi and Parisi, 1991].

5. Identifying optimal parameter values requires large problem spaces to be searched. This is computationally expensive [deGaris, 1995, Fleischer, 1995, Gruau, 1994] and may ultimately be intractable [Dellaert and Beer, 1996]. Parameter searches primarily use function as the fitness measure and are generally initiated without any pre-conception of the DANN's final structure.

3.2 A DANN to Investigate Self-Organisation

The DANN presented in this document incorporates the desirable features of existing DANNs, whilst aiming to achieve a balance between the level of biological detail modelled and the ease with which the model can be manipulated.

A minimal set of developmental rules is used that is intentionally simple. The rules are inspired by the belief that biological neural development uses simple rules [Stryker, 1994] and by Reynolds' *boids* rules [Reynolds, 1987]. The key feature of the rules are that they are an abstraction of biological self-organisation. To explicitly model biological development rules requires detailed knowledge of molecular and chemical processes, that vastly increases the complexity of modelling. Using simple, locally-applied rules vastly reduces the computational load of simulations and enables modelling of large-scale systems. A restricted number of rules is used such that the size of the problem space is minimised when searching for optimal parameters.

Development can occur in simulated 2D and 3D chemical environments. Neurons are modelled as discrete elements which locally interact with each other and the environment. The simulated environment allows more complex cell-cell interactions than previously proposed [Nolfi and Parisi, 1991, Vaario, 1994] but is less complex than that modelled by Fleischer [Fleischer, 1995].

4 A Model of Biological Self-Organisation

4.1 Starting Assumptions

The model of biological self-organisation proposed is based on the following assumptions:

1. The development process occurs within a simulated chemical environment. This model is intended to be a simplified version of the environment in which biological development occurs (see 2.1.1).

2. Cell division, differentiation and migration have previously occurred and resulted in layers of neurons as found in the retina and visual tract [Van Essen and Anderson, 1995].

3. The distribution of neurons within layers is inherently regular.

4. Development is determined by simple, local rules. As stated previously these rules are abstract versions of those governing biological development.

4.2 The Three Phases of Biological Self-Organisation

Having established the starting conditions, the development process can be considered as three phases (see Figure 3). Each phase is a specific set of self-organising mechanisms and processes.

1. **Phase 1 - Growth.** The creation of inter-layer and inter-neuron connections.

2. **Phase 2 - Spontaneous Activity.** Refinement of the initial skeletal structure by spontaneous activity.

3. **Phase 3 - Learning.** Further refinement of structure and initial pattern learning based upon external stimuli.

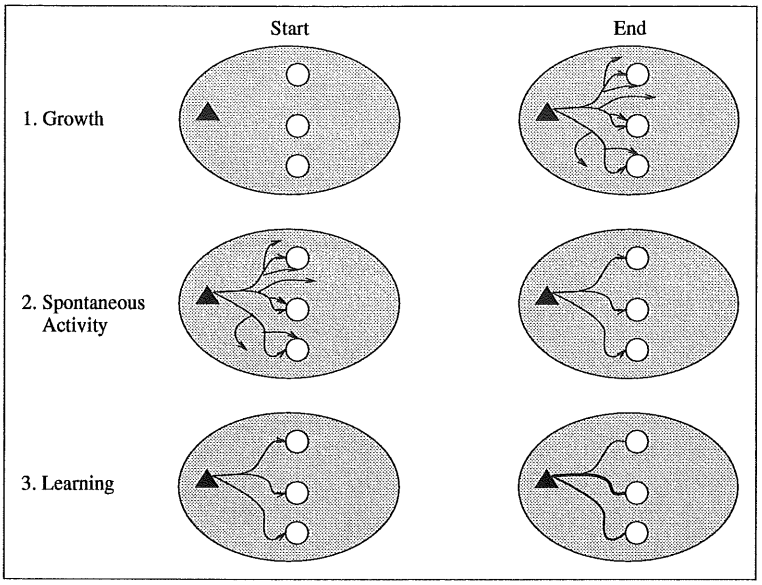

Figure 3: The three phases of biological self-organisation.

The specification of the model enables it to be applied at different levels of self-organisation as shown in Figure 1. Individual components may either be neurons, leading to the development of neuron-neuron connectivity, or networks of neurons, creating modular systems.

Although there is a distinct transition between the spontaneous activity and learning phases [Goodman and Shatz, 1993], the three phases of biological self-organisation overlap to varying degrees. In biological systems for example the growth phase and the spontaneous activity phase are interdependent [Shatz, 1994]. An interpretation of the overlap and significance of the three stages is shown in Figure 4.

This definition of self-organisation emphasises the relationship between function and structure [Rust and Bolouri, 1996]. Underlying neural structure places boundary conditions and restrictions on learning and the resulting functionality.

4.3 Methodology

Function and structure are intertwined in neural systems. In addition to the topological interconnection pattern between neurons, the relative location of synapses on dendritic trees plays a crucial role on the effect of the corresponding inputs on post-synaptic activity. This is true not only for combinations of excitatory and inhibitory synapses [Koch et al., 1983] but also for groups of excitatory only synapses [Mel, 1994]. Since it is possible to produce different networks with the same function using many different developmental routines, it is unknown which developmental rules produce the best building blocks for artificial evolution. We aim to search for developmental rules which produce structures, not functions, which mimic biological neural networks. By basing our studies on the retina, the need to learn functionality is avoided and a greater emphasis is placed on the relationship between function and structure [Rust and Bolouri, 1996].

It is believed that if the 'right' architecture can be grown then function will follow. By examining biological neural systems 'target' structures can be defined which can be used to measure the fitness of grown structures. Having a single desired architecture constrains and directs the search for optimal developmental parameters thereby reducing development time and computational cost. It is hoped that there will be no unique set of developmental parameters for each target structure. If the developmental growth rules are robust and adaptable then multiple sets of parameter values will result in closely matching structures.

The retina is used as the 'target' structure in this document, since it has a stereo-typed architecture [Van Essen and Anderson, 1995] and may be assumed to be reasonably 'hard-wired'. Hence since learning (synapse modification) is not considered to be important in the retina, this allows the experimental work

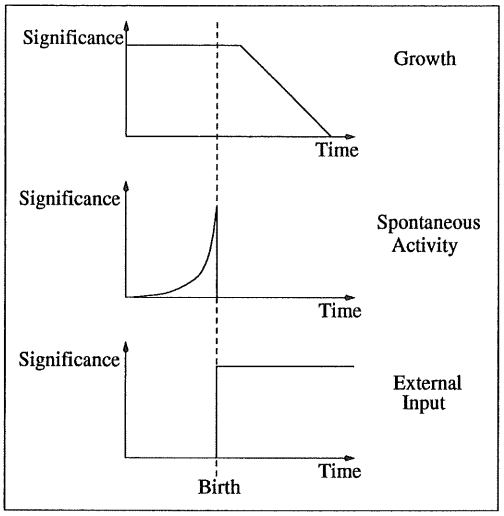

Figure 4: Overlap of the three phases of biological self-organisation.

to concentrate on the first two phases of biological self-organisation. Once the developmental rules and parameters have been characterised for the retina, learning can be incorporated when modelling higher levels in the visual cortex.

The remainder of the document discusses the growth of neuron-neuron connections: Phase 1 of *the three phases of biological self-organisation* model applied at the level of neuron connectivity (see Figure 1).

5 Phase 1: The Modelling Environment

5.1 Chemical Environment

Networks develop within a virtual chemical environment, which can be either modelled as a plane (2D) or as a box (3D). An underlying gradient is imposed on the substrate of the environment to encourage neurites to move in predefined directions. The substrate chemical gradient is maximum in one direction and zero in perpendicular directions. The change in gradient is small compared to the chemical gradients emitted by the neurites (see 5.2). Axon growing neurons are positioned at the bottom of the gradient whilst dendrite growing neurons are found at the top of the gradient.

5.2 Neuron Model

Neurons are fixed within the environment and contain a set of internal genetic rules which dictate growth. The rules enable neurons to extend both axons and dendrites (neurites) though not simultaneously. The same rules are used for both axon and dendrite processes.

Neurites emit chemicals from every point along their length, which produce local gradients in the environment. This process is inspired by the production of neurotrophins in biological systems (see 2.1.1). The gradient fronts are spherical in 3D and circular in 2D. Axon producing neurons emit negative chemicals whilst dendritic neurons emit positive chemicals. The characteristics of a neurite's emitted gradient are determined by the parameters; strength, range and diffusion law ($1/\text{distance}^n$). Gradients are assumed to be additive.

The tips of extending neurites actively sense the local chemical environment, analogous to biological growth cones. The internal rules determine the response of the neurite: extend, split(branch) or terminate if a target neurite is contacted. Hence although the behaviour of a neurite is determined globally at the neuron level, the response is local to the tip of the neurite. Neurites are restricted in the directions of movement: in 2D a neurite can move in one of 8 directions and in one of 26 in 3D.

5.3 Growth Rules

Neurite growth is governed by the following simple developmental rules:

1. Follow the path of the steepest gradient of the opposite polarity to that emitted. Therefore, axons perform hill-climbing and dendrites gradient-descent.

2. Maintain the same path unless attracted by a larger gradient. This allows neurites to be guided initially by the underlying substrate gradient and then to search out local target neurites over time.

3. Split if genetically programmed or if there are two strong local gradients.

It is proposed that no one rule will account for all neurite growth and that it is a combination of these rules which will generate the final, complex structures.

6 Phase 1: Implementation

6.1 Growth Initialisation

Neurons are arranged into layers separated by pre-defined distances. All neurons in the same layer possess the same set of genetic rules but the range and strength of the chemical emitted by each neuron may vary.

6.2 Growth Sequence

Growth occurs between two layers at a time forming connections when axons and dendrites meet. An extension to this process allows connections made in a previous growth sequence to become the targets for neurites from another layer; the connections become static targets and do not grow neurites themselves. Growth is halted after a pre-defined period, after which the resulting chemical gradients are removed and the underlying substrate gradient replaced, before subsequent layers are processed.

Although this is much simplified in terms of biological development, each layer may be considered as one class of neuron which is sensitive to the chemical labelling of only one other class of neuron. Layers can therefore be considered in pairs, in isolation from the rest of the system.

6.3 Neurite Splitting

A neurite may split into two separate paths moderated by two rules. Environmental splitting is a stochastic process regulated by the local chemical gradients, whilst genetic splitting is deterministic being regulated by the neuron itself. This mixture of deterministic and non-deterministic behaviour is believed to have greater flexibility and to be capable of generating more complex structures.

Splitting is not a limitless procedure. A *recovery period* is incorporated such that once a neurite has split it is inhibited from splitting for a predetermined time. It is proposed that a neurite expends resources in splitting which must be replaced before splitting can occur again. Split inhibition prevents a neurite splitting profusely once it is close to a target neurite. Upon splitting the strength of the chemical gradient emitted is halved, whilst the range and diffusion law remain the same.

6.3.1 Environmental Splitting.

Environmental splitting uses the two largest local gradients to calculate a probability of splitting. Neurites are more likely to split if any 2 gradients are large and approximately equal. The split occurs in the directions of the two largest gradients. When split inhibition is enforced it will prevent an environmental split from occurring. Hence although environmental information is the stimulus for splitting, environmental branching is regulated by the internal genetic rules.

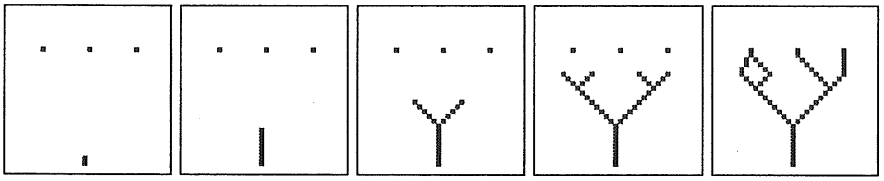

Figure 5: An example of 2D neuron growth: Sequential steps of development.

6.3.2 Genetic Splitting.

Unlike environmental splitting this branching occurs at pre-determined time intervals. Intervals were chosen such that multiple branching occurs at early stages of growth reducing in frequency over time.

As with environmental splitting, there is an interaction between the genetic split signal and the direction of split. Once a genetic split has been signalled the direction of the split is determined by the gradients at the tips of neurites. A number of genetic splitting methods have been implemented. For example, a split may occur either side of the maximum gradient or either side of the current direction of movement. (A full description of splits may be found in [Rust et al., 1996].) In 3D the pair of directions branched along is chosen at random with the branches being in the same plane.

6.4 Connection Formation

When a growth cone collides with either a trail, tip or cell body of a target neuron a connection is formed. The neurite which makes contact stops growing. Further connections can be made on existing connections.

7 Phase 1: Results

7.1 Simple Examples

Figure 5 shows a 2D example of neuron growth. Three static neurons are the targets for a single, axon growing neuron. The target neurons are equally spaced whilst the growing neuron is offset to the left of the central target neuron. The growing neurites maintain their current paths over the underlying substrate gradient until attracted by the local gradients of the target cells. The neurite splits are genetically determined and are either side of the local gradient.

Figure 6 is an example of a 3D neuron system containing two dendrite growing neurons and one axon emitting neuron. Cell bodies are represented as spheres. This example highlights some of the primitive structures which can be grown using the developmental rules.

7.2 A 3D Retina Model

The retina structures grown so far are modelled on the cells that make connections in the outer plexiform layer: cones, horizontal and bipolar cells.

An example of a 3D model of a 16 neuron retina is shown in Figure 7, where cone cells are the lower layer, horizontal cells centre and bipolar cells uppermost.

3D structures are visualised using VRML (Virtual Reality Modeling Language). Examples can be found at the following website: *http://erdc.herts.ac.uk/NAGweb/alistairR/vrml/vrml.html.*

8 Discussion

The results presented are for the first phase of the proposed model of biological self-organisation: neural growth. This was investigated at the level of neuron-neuron connectivity. The DANN used is believed to be the first documented 3D implementation. Alternative 3D environments have only been suggested [Fleischer, 1995]. Although the model implemented is not argued to be biologically plausible, the resulting

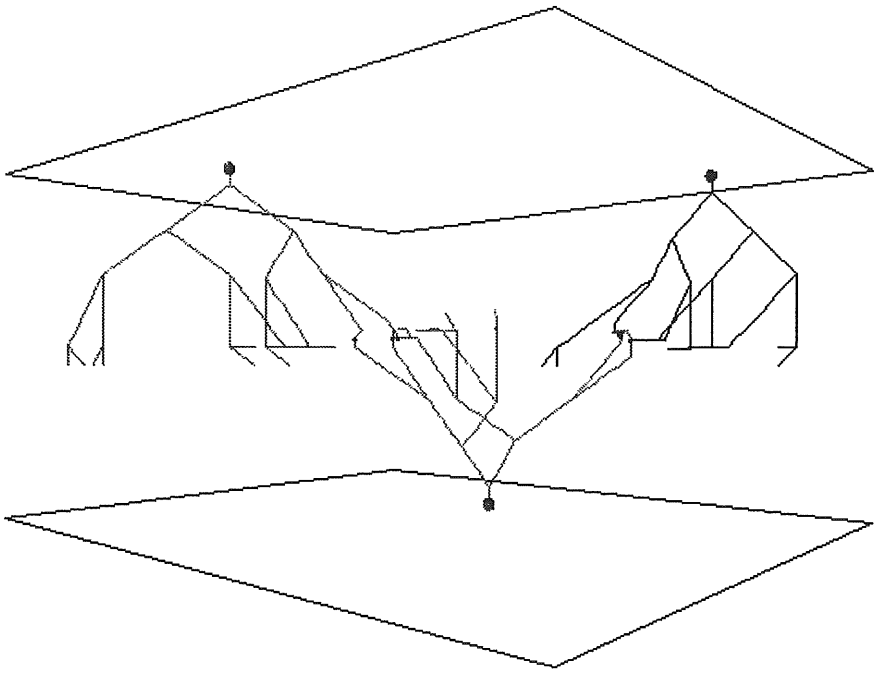

Figure 6: An example of 3D neuron growth.

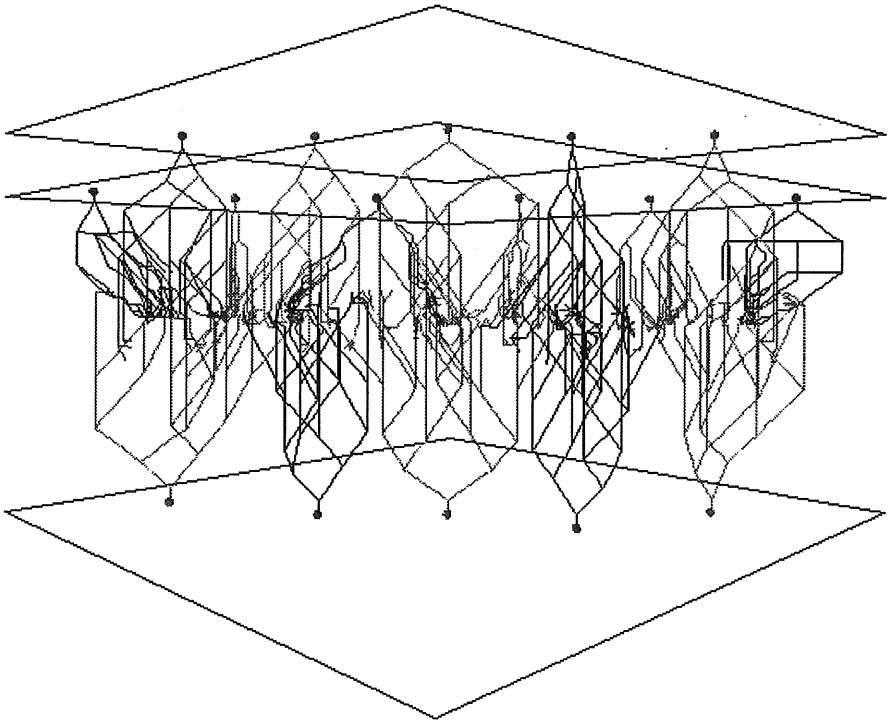

Figure 7: A 3D example of a model retina.

structures are neural-like in appearance. This is due both to the developmental rules used and the simulated chemical environment.

The developmental rules were specifically designed to be simple and to be applied locally. The rules however have been shown to be capable of creating complex, multi-layer neural structures. It is thought that because of the simplicity of the rules they are more flexible than more rigidly defined mathematical encodings. The mixture of deterministic (genetic) and non-deterministic (environmental) splitting rules also permits more complex neurite interactions. Allowing growth onto existing connections further increases structural complexity and may be required at subsequent stages of self-organisation.

The simulated chemical environment allows greater interactions between the neurons and neurites than most previous approaches. It is felt that this environment, although representing a simplified model of biology, exhibits the principle mechanisms required for modelling self-organisation. The current description of the model contains a sufficient degree of biological detail without being computationally expensive and inflexible compared to more complex models [Fleischer, 1995].

The simultaneous growth of axons and dendrites is also novel. A number of DANNs implement neurite growth but this is either simplified by static target cells [Jakobi, 1995, Vaario, 1994] or growth is unaffected by neighbouring cells [Nolfi and Parisi, 1991]. The growth of two classes of neurite allows a greater range of cell-cell interactions.

The methodology proposed represents a unique use of a DANN. The modelling is directed at mimicking biological structures using biologically inspired rules. Other DANN models use ad-hoc developmental rules to evolve artificial networks for specific applications. By closely modelling biological structure we pose the question: can such simple, abstract rules provide a description of the way biological neural structures develop in nature ? If simple, computational rules can be found and optimally encoded, this provides the basis on which artificial evolution can be efficiently investigated.

The results presented so far demonstrate that many more connections between individual neurons are made than may be necessary. Hence the primary role of phase two (refinement using spontaneous activity) must be to prune extraneous initial connections. The next stage of the study will investigate the roles of spontaneous activity and neural regulation in refining connections in phase 2 [Alberts et al., 1994, Purves and Lichtman, 1985, Shatz, 1994].

Biological structures have been used to direct the growth of the model retina towards some 'target' morphology. However, a more exacting structural definition is required than the general model currently used. The emphasis on the relationship between structure and function dictates this requirement since it is proposed that function will follow from the correct 'target' structure.

The developmental parameters used thus far have been chosen based on visual inspection of the structures grown. A genetic algorithm is being used to investigate optimal development parameter values. It is hoped that the search will demonstrate the robustness of the developmental rules. Exploring the parameter space of the developmental rules will determine the significance and sensitivity of parameter values on resulting structures.

9 Conclusion

This document has presented preliminary work involved in the design of a shape recognition system inspired by biological visual systems. The study sits within a larger programme of work aimed at investigating the emergent self-organisation of structure that occurs during neural development. It is suggested that biologically inspired developmental models provide an appropriate starting point for evolutionary design of sophisticated neural systems.

A theoretical framework has been proposed which incorporates both a description of self-organisation across scales of structure and the grouping of self-organising processes into three phases. This emphasis on structure enables biological neural architectures to be defined as *target* structures, directing the development procedure. It is believed that function will follow from having the 'right' structure.

A bottom-up approach has been adopted starting with the investigation of neuron-neuron connectivity growth. The developmental growth rules used are inherently simple, representing an abstraction of the key neural development processes. Growth occurs within a DANN testbed, which can simulate both a 2D or 3D chemical environment. A general three layer structure of part of a retina has been developed in 3D, demonstrating how simple, self-organising rules can develop complex neural-like structures.

References

[Alberts et al., 1994] Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J. D. (1994). *Molecular Biology of the Cell.* Garland Publishing, London, 3rd edition.

[Boers et al., 1993] Boers, E. J. W., Kuiper, H., Happel, B. L. M., and Sprinkhuizen-Kuyper, I. G. (1993). Designing modular artificial neural networks. Technical Report 93-24, Leiden University.

[deGaris, 1995] deGaris, H. (1995). CAM-BRAIN, the evolutionary engineering of a billion neuron artifical brain by 2001 which grow/evolves at electronic speeds inside a cellular automata machine (CAM). In Pearson, D. W., Steele, N. C., and Albrecht, R. F., editors, *Proceedings of the 2nd International Conference on Artificial Neual Networks and Genetic Algorithms*, pages 84–87, New York. Springer-Verlag.

[Dellaert and Beer, 1994] Dellaert, F. and Beer, R. D. (1994). Co-evolving body and brain in autonomous agents using a developmental model. Technical Report CES-94-16, Case Western Reserve University.

[Dellaert and Beer, 1996] Dellaert, F. and Beer, R. D. (1996). A developmental model for the evolution of complete autonomous agents. In *To be presented at the 4th International Conference on Simulation of Adaptive Behaviour (SAB96).*

[Dowling, 1992] Dowling, J. E. (1992). *Neurons and Networks : An Introduction to Neuroscience.* Harvard University Press, London, 1st edition.

[Fleischer, 1995] Fleischer, K. (1995). *A Multiple-Mechanism Developmental Model for Defining Self-Organizing Geometric Structures.* PhD dissertation, California Institute of Technology.

[Goodman and Shatz, 1993] Goodman, C. S. and Shatz, C. J. (1993). Developmental mechanisims that generate precise patterns of neuronal connectivity. *Cell*, 72(Review Suppl):77–98.

[Gruau, 1994] Gruau, F. (1994). *Neural Network Synthesis Using Cellular Encoding and the Genetic Algorithm.* Thesis, Ecole Normale Superieure de Lyon.

[Jakobi, 1995] Jakobi, N. (1995). Harnessing morphogenesis. In *To Appear in Proceedings of International Conference on Information Processing in Cells and Tissues.*

[Kitano, 1990] Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4:461–476.

[Koch et al., 1983] Koch, C., Poggio, T., and Torre, V. (1983). Nonlinear interactions in a dendritic tree: Localization, timing and role in information processing. *Proceedings of the National Academy of Sciences of the USA*, 80(9):2799–2802.

[Mel, 1994] Mel, B. W. (1994). Information-processing in dendritic trees. *Neural Computation*, 6(6):1031–1085.

[Nolfi and Parisi, 1991] Nolfi, S. and Parisi, D. (1991). Growing neural networks. Technical Report PCIA-91-15, Instituite of Pyschology, Rome.

[Purves and Lichtman, 1985] Purves, D. and Lichtman, J. W. (1985). *Principles of Neural Development.* Sinauer Associates, Sunderland, MA, 1st edition.

[Reynolds, 1987] Reynolds, C. W. (1987). Flocks, herds, and schools : A distributed behavioral model. *Computer Graphics*, 21:25–34.

[Rust and Bolouri, 1996] Rust, A. G. and Bolouri, H. (1996). Developmental neural networks for shape recognition: Motivation and review. Technical Memorandum ERDC/1996/0006, ERDC, University of Hertfordshire, UK.

[Rust et al., 1996] Rust, A. G., George, S., Bolouri, H., and Adams, R. (1996). Developmental artificial neural networks for shape recognition: A model of the retina. Technical Memorandum ERDC/1996/0011, ERDC, University of Hertfordshire, UK.

[Shatz, 1994] Shatz, C. J. (1994). Role for spontaneous neural activity in the patterning of connections between retina and LGN during visual system development. *International Journal of Developmental Neuroscience*, 12(6):531–546.

[Stryker, 1994] Stryker, M. P. (1994). Precise development from imprecise rules. *Science*, 263:1244–1245.

[Vaario, 1994] Vaario, J. (1994). Modeling adaptative self-organization. In Brooks, R. A. and Maes, P., editors, *Artificial Life IV : Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 313–318, London. MIT Press.

[Van Essen and Anderson, 1995] Van Essen, D. C. and Anderson, C. H. (1995). Information processing strategies and pathways in the primate visual system. In Zornetzer, S. F., Davis, J. L., Lau, C., and McKenna, T., editors, *An Introduction to Neural and Electronic Networks*, pages 45–76. Academic Press, London, 2nd edition.

[von der Malsburg, 1995] von der Malsburg, C. (1995). Network self-organization in the ontogenesis of the mammalian visual system. In Zornetzer, S. F., Davis, J. L., and Lau, C., editors, *An Introduction to Neural and Electronic Networks*, pages 447–463. Academic Press, London, 2nd edition.