Evolutionary Neural Topiary†: Growing and Sculpting Artificial Neurons to Order

Alistair G Rust[1,2], Rod Adams[1] and Hamid Bolouri[2,3]

[1]Department of Computer Science, University of Hertfordshire, UK
[2]Science & Technology Research Centre, University of Hertfordshire, UK
[3]Division of Biology, California Institute of Technology, USA
http://strc.herts.ac.uk/NSGweb/ {*a.g.rust, r.g.adams, h.bolouri*}*@herts.ac.uk*

Abstract

Designing artificial systems with ever more biologically-plausible 'brains' continues apace and permits investigations into the computational capabilities of engineered systems. Creating artificial neurons with biologically-realistic morphologies is however a non-trivial problem. This paper addresses *growing neurons to order*, neurons with morphologies exhibiting strong biological traits. A biologically-inspired simulator of neural development is coupled with a genetic algorithm to evolve 3-dimensional neuron morphologies. The morphology of a biological neuron provides the exemplar target against which the developmental evolution process is gauged.

Realising the Potential of *Brain Building*

During the infancy of Artificial Life, the goal of building artificial organisms and systems with artificial brains – *Brain Building* (deGaris, 1990) – was often a mooted topic. As the domain of Artificial Life has matured, so has *Brain Building*, such that a number of leading-edge, research corporations currently invest substantial resources into developing artificial brains (RIKEN, 2000; NASA, 2000; ATR-HIP, 2000). Much of the maturation has come through the constant increase in computational resources which provide the means with which to simulate artificial neural systems to ever higher levels of biological-plausibility (Hines and Carnevale, 1997).

With biological-plausibility comes the opportunity to fully exploit the complexities of neural computation. Single neurons, let alone networks and systems, possess adaptive, dynamic computational capabilities (Koch and Segev, 1998). One of the key determinants of these capabilities is the relationship between a neuron's morphology and its function (Mainen and Sejnowski, 1996). However, only a handful of computer simulations have explored the relationship between function and form in biologically-plausible terms, e.g. (Mel, 1994).

Our research explores the feasibility of evolving developmental programmes that create biologically-plausible structured neural systems (Rust, 1998). Previously we have used simulated neural development to grow artificial neurons that were functionally evaluated against biological neurons (Rust and Adams, 1999). Experiments showed a strong relationship between function and form in artificial neurons where for example, implausible morphologies possessed inappropriate functionality. Even the final neurons, although functionally similar, had morphologies which still differed from their biological targets. So to mimic the computational capabilities of real neurons and to speed up the search for artificial equivalents, we argue that a method of creating biologically-cogent artificial neurons is required. Namely we are interested in *growing neurons to order*. In this paper the evolution of neuron morphology is explored.

Modelling Neural Development

Modelling of neural morphology using the re-writing rules of L-systems (Lindemayer, 1968) has been explored by a number of groups (Burton et al., 1999; Ascoli, 1999). (Ascoli, 1999) in particular has been using L-system rules whose grammar is derived from observations of biological neural morphology (e.g. typical branching angles and rate of dendritic diameter reduction at branch points). L-system neuron models however, do not typically allow interactions between a growing neuron and its environment. The development of a neuron and hence its visual form are dependent on the implementation of the hand-crafted re-writing rules.

Fleischer and Barr developed an extensive simulator which incorporated many self-organising bio-physical phenomena expressed as differential equations (Fleischer, 1995). Attempts were made at evolving neural morphology but these were computationally intensive due to the large parameter search space involved.

We have implemented a 3D model of neuro-biological development, in which neuron-to-neuron connectivity is created through interactive self-organisation (Rust et al., 1998; Rust, 1998). Development occurs as a number of overlapping stages, which govern how neurons extend axons and dendrites, collectively termed neurites. Neurons grow within an artificial embryonic environment, into which neurons and their neurites emit local chemical gradients. The growth of neurites is influenced by

†*topiary*: a branch of gardening, the clipping of trees into imitative and fantastic shapes.

the local gradients and the sets of interacting, developmental rules. The *interactive rules* enable neurites to navigate and branch in response to local developmental conditions, and to prune unwanted connections.

The developmental rules are controlled by parameters, much in the same way as gene expression levels can be thought of as parameters for biological development. By varying these parameters, a variety of neuron and network morphologies can be achieved. Examples of individual neurons are illustrated in Figure1.

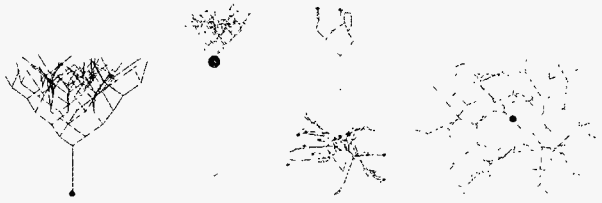

Figure 1: Examples of 3 dimensional neurons grown using the developmental simulator.

Evolving the developmental model for a specific network then becomes equivalent to the search for optimal sets of developmental parameters using, for example a genetic algorithm (GA). Previously we have used a GA to evolve developmental parameters which lead to the creation of an edge-detecting retina (Rust et al., 1998). In this paper the simulator is used to grow single neurons within developmental environments rich with gradients of chemical attractants. The single neurons are then evaluated in terms of their morphological characteristics against the dendritic structure of a biological neuron.

Analysing Neural Morphology

Numerical Evaluation

In order to verify the match between an artificial neuron and a desired biological counterpart, means of characterising biological neurons are required. However, no one set of benchmark measures (quantitative, topological and/or qualitative) exists within neuroscience literature, with which separate classes of neurons can be reliably characterised or compared. Hence in the experiments reported in this paper, as with complementary work (Ascoli, 1999), a minimal set of characteristics is specifically selected. The chosen set aims to reliably represent the geometrical and spatial characteristics of the target biological neuron in the least terms.

Some approaches seek to find generic solutions for classes of neurons by averaging measures from neuron databases. In this paper we aim to gain an understanding of how the morphological characteristics of individual neurons arise before attempting more generic, class-based approaches. Consequently we seek a methodology to clone artificial neurons from a given target neuron. The chosen exemplar neuron for this paper, a layer 5 *pyramidal* neuron, is shown in Figure 2.

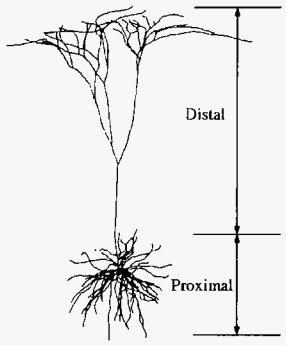

Figure 2: The dendritic morphology of the target, layer 5 *pyramidal* neuron, where the cell soma is located at the centre of the proximal dendritic tree. Morphology obtained from (Mainen and Sejnowski, 1996).

Visual Evaluation

One proviso in using only representative numerical characteristics is, that although particular artificial solutions may satisfy the selected criteria, they may not necessarily be solutions which are visually satisfying. For example, although a neuron may be evolved to have the correct number of tips (terminal segments), it does not mean that the overall geometry of a solution is necessarily consistent with its biological counterpart. In this paper we therefore examine evolving artificial neurons to criteria based upon numerical characteristics alone as well as criteria based on visual appearance alone.

Simply using a GA to traverse the parameter space looking for visually desirable morphologies is not however an effective method. In early generations due to the large diversity in initial populations, visual discrimination will be slow since most morphologies will be discarded until relevant, morphological search sub-spaces are found. A more directed search, where the paths between generations of morphologies could be more directly traversed, would be desirable in this instance.

One such directed approach was used by Dawkins to evolve his artificial *biomorphs* (Dawkins, 1991). An analogous approach towards the evolution of neuron morphology based on visual discrimination is used here. Developmental parameters are selected to form a genome. Each generation of neuron is created by selecting each gene in turn and creating 2 mutated copies of the gene by simply incrementing and decrementing its current value by a pre-determined delta. The values of the remaining genes in the genome remained fixed. Genes are given minimum and maximum values which they can not exceed and a start value is randomly selected from within this range. For each generation, a neuron for each genome is grown and displayed on the screen. Based on a visual comparison between the evolved neuron morphologies and the target biological neuron, a morphology is selected for the next generation.

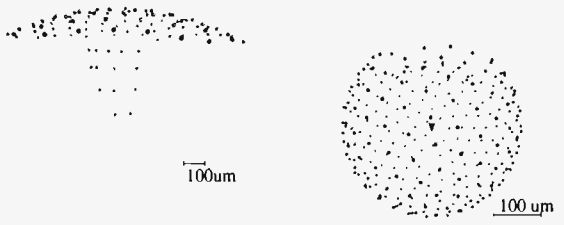

(a) Distal (b) Proximal

Figure 3: Arrangement of chemical attractants for the 2 developmental environments. The soma of the pyramidal neuron is the lowest object in (a) and the cone at the centre of the sphere of attractants in (b).

Developmental Strategies

Self-Organising Interactions

A feature of the developmental simulator is the ability to influence or sculpt the growth of neurons by exploiting the potential of the interactive, self-organising mechanisms. This was achieved by using the placement and temporal expressions of chemical attractants to guide dendritic growth. For example, the placement of attractants can allow asymmetrical growth, orientated along one particular axis. Being able to specify characteristics of the attractants, provides a more intuitive approach to growing neurons to order. This can avoid relying upon adapting developmental rules and parameters, which can often feel like tinkering with a *black box*.

Developmental environments of attractants were therefore constructed to interact and sculpt the neurons as they grow. These are illustrated in Figure 3. In the case of the proximal dendritic tree, its morphology is approximately radial and symmetrical, hence an environment of attractants placed on the surface of a sphere is used (Figure 3(b)). The dendritic tree is not however perfectly symmetrical as the upper dendrites have shorter lengths compared to their neighbours. To sculpt the artificial neurons to this subtle degree of asymmetry, a small proportion of the attractants are positioned marginally closer to the centre of the sphere.

The morphology of the distal dendrites is directly reflected in the arrangement of attractants (see Figure 3(a)). Intermediate attractants are placed to guide dendrites vertically before the arc of attractants at the top of the environment are encountered. In both sets of developmental environments the attractants were placed to reflect the spatial dimensions (μm) occupied by the biological neuron.

Phased Development

Neurons are known to respond to different environmental cues in accordance with their spatial and temporal locations (Hall, 1992). Trying to reproduce the shape of a complex dendritic tree is therefore unlikely to succeed using an algorithmic model with global settings.

One potential solution to this problem is to split the development of the morphology into different phases and optimise these phases separately. Criteria for subdividing the morphology can then be based on such properties as branching frequency and the orientation of growth. For each phase of development growth control parameters can be given different ranges of potential values at different times. Branching for example, may then be controlled by inhibitory values when it is not required and excitatory values when it is required.

The development of the artificial pyramidal neuron is divided into 3 such phases. The first phase consists of the initial vertical development of the dendritic tree away from the soma, where preliminary branches are established. The effects of attractants are time dependent during this phase to induce interactive branching at the desired times. The end points of phase 1 are then used as the starting points for the second phase of branching, directed towards the arc of uppermost attractants. The final phase produces the development of the proximal dendrites into the environment of attractants arranged on the surface of a sphere (see Figure 3(b)).

A complete pyramidal neuron is achieved by combining all 3 phases of the developmental process.

Results

Since the proximal dendrites are basically 'spherical', visual selection is neither useful nor necessary. Instead, we used automated evolution based on numerical measures of morphology. Visual selection was used to sculpt the asymmetrical distal dendritic tree.

Visual Selection: Distal Dendrites

For both phases 1 and 2, 4 developmental parameters were evolved. Any number of parameters could have been chosen but 4 parameters which have key effects on branching interactions were selected. Specifically the parameters adapted were: (i) the probability of branching as a function of local attractant gradients, (ii) inhibition of branching following a previous branch, (iii) branch inhibition due to saturation caused by nearby attractants and (iv) local repulsion between growing dendrites.

Figure 4 shows the various stages of the visual selection process. Since the genome of the neuron contained 4 genes, 8 morphologies are grown for each generation (each gene is mutated both positively and negatively whilst the other genes remain fixed). A typical screenshot of an intermediate generation of phase 1 is shown in Figure 4(a). The finally selected morphology is illustrated in Figure 4(b). The morphology for phase 2 was evolved in the same manner.

Numerical Selection: Proximal Dendrites

Six morphological properties of the proximal dendrites were chosen to be the criteria upon which evolved neu-

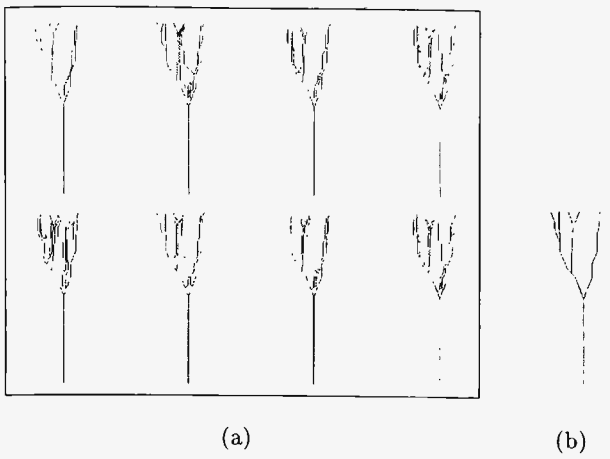

(a) (b)

Figure 4: Typical images from phase 1 of the visual selection process. (a) A screen-shot of one generation of 8 neurons during an intermediate stage (25 generations). (b) The final, visually selected neuron (52 generations).

rons were evaluated. The criteria were: (i) total length of the dendritic tree, (ii) number of dendritic segments (i.e. the number of dendritic lengths between branch points (iii) number of tips, (iv) average length of tips, (v) length of the longest tip and (vi) the centre of gravity. (vi) was aimed at providing a spatial perspective.

These properties were extracted from data from the neuron and were chosen to effectively describe the morphology of the neuron in the least possible terms. The error value of a neuron was calculated using:

$$e = \sum_{i=1}^{6} \frac{|a_i - t_i|}{t_i} \qquad (1)$$

where a_i is the value of a property of the evolved neuron and t_i is the target value for the proximal dendrites.

The GA used in these experiments was GENESIS (Grefenstette, 1990). 14 developmental parameters were encoded in the GA using 44 bits. The encoded parameters controlled the times at which dendrites could branch and how the growing tips would interact with the attractants in its environment. The population size was 50 and each population was randomly initialised. The crossover rate was 0.6 using dual point crossover. The mutation rate was set such that at each generation approximately 15% of the population would undergo a bit mutation. The role of the GA was to minimise the error from the evaluation function (1).

Ten simulations were performed on a 400MHz Pentium PC running Linux. Each population of neurons was evolved for 150 generations where each population, on average, took 50 minutes to grow and to be evaluated. The best individual neuron evolved had a fitness value of 0.195, which is equivalent to an average error of 3.25% per evaluated property. The average error for the 10 experiments was 0.296 (4.93% per property).

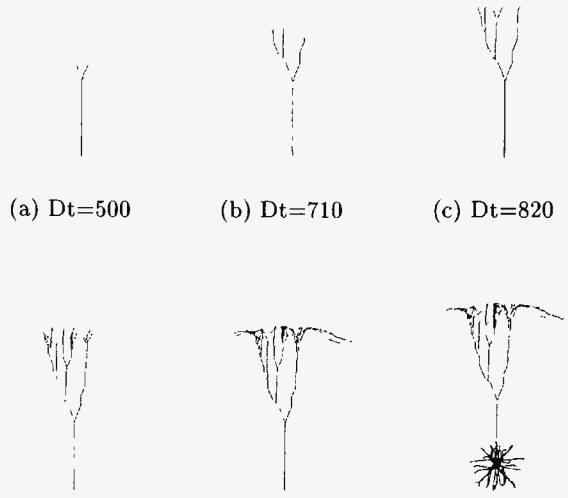

(a) Dt=500 (b) Dt=710 (c) Dt=820

(d) Dt=1000 (e) Dt=1620 (f) Dt=1820

Figure 5: Snapshots of the growth of the pyramidal neuron in all 3 developmental phases, at various developmental times (Dt). (c) is the end of phase 1, (e) is the end of phase 2 and (f) shows the final phase.

Combining the Phases

Figure 5 shows the complete developmental process combining all 3 phases of the developmental evolution process. The morphology of the evolved artificial neuron is compared against its biological target in Figure 6.

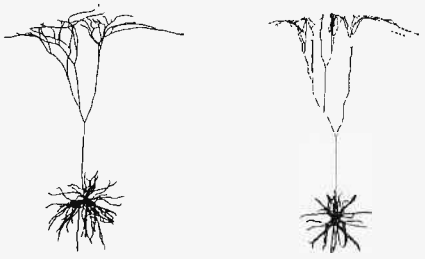

(a) (b)

Figure 6: A comparison of morphologies between (a) the target biological neuron and (b) the best evolved neuron.

Discussion

The evolution of the proximal dendrites demonstrated that if visual discrimination of a neuron is not critical and its dendritic tree can be assumed to be approximately symmetrical, then numerical measures alone can be sufficient. Due to the radial, symmetrical morphology of the proximal dendrites a regular spacing of attractants

was adequate. Growing neurons in this way can then be a blind, automatic process.

Where the morphology of the dendritic tree is more complex, then a blind process may fail. To evolve the morphology of the distal dendrites using numerical selection alone, would require a larger set of numerical characteristics to be extracted and used as guides for development. Choosing representative measures is a non-trivial task and there is no guarantee that these measures can be effectively interpreted by the evolutionary process to produce desired morphologies. (Evolutionary computation (EC) literature contains many examples of algorithms adapting to unwanted behaviours and functions.) Evolving using visual examples of morphologies enables a more directed approach, somewhat more intuitive than a *black box* method. Greater control of the developmental process was also afforded by splitting the morphology of the pyramidal neuron into separate, tractable sub-problems.

The modelled mechanisms of interaction encapsulate the self-organising principles inherent in bio-physical processes. These interactions are shared between a growing neuron entity and its developmental environment. Due to these interactions, fewer parameters which control the growing behaviour of the neuron need to be encoded and evolved in the genome. This thereby leads to a reduction in the size of the genome under evolution, which has 2 potential benefits. Firstly, it reduces the reliance on the chosen EC tool alone to identify optimal solutions by modifying the developmental control parameters. The *evolutionary workload* can be more evenly divided between the bio-physical interactions of the developmental model and the EC method. Secondly, reduced genome sizes should lead to shorter evolutionary searches.

Future Work

As well as validating the model further on other exemplar neurons, the coupling of evaluation of evolved morphologies with dynamical membrane models needs to be explored. Previous work (Rust and Adams, 1999) explored the relationship between function and morphology but without directing the growth process to produce more visually biological-like morphologies. The question to be addressed is, having selected a morphology based on its similarity to a biological neuron, does functionality (e.g. the pattern of spike trains) come for 'free'?

During previous experiments evolving artificial neurons with functional characteristics akin to biological examples (Rust and Adams, 1999), the major bottleneck was the computational cost of evaluating evolved neurons with the chosen compartmental membrane model. On-going work aims to identify a computationally less intensive membrane model to permit faster evaluations.

Conclusions

A biologically-inspired developmental simulator subject to evolutionary adaptations, was presented in this paper which was capable of generating biologically-realistic neuron morphologies. Growing dendritic trees were sculpted through self-organising interaction with chemical attractants positioned within the developmental environment. In-conjunction with this, dividing the search for the morphology of a complex neuron into separate and solvable phases proved to be a beneficial strategy. Such a combination of techniques and strategies provides a jumping off point to explore even more dynamic artificial neurons and networks.

References

Ascoli, G. A. (1999). Progress and perspectives in computational anatomy. *Anatomical Record*, 257(6):195–207.

ATR-HIP (2000). ATR, Human Information Processing Laboratories. `www.hip.atr.co.jp/`.

Burton, B. P., Chow, T. S., Duchowski, A. T., Koh, W., and McCormick, B. H. (1999). Exploring the brain forest. *Neurocomputing*, 26-27:971–980.

Dawkins, R. (1991). *The Blind Watchmaker*. Penguin, London, 1st edition.

deGaris, H. (1990). Building brains with GenNets. *Proc. of the Int. Neural Network Conf.*, pages 1036–1039, Dordrecht, Netherlands. Kluwer Academic Publishers.

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

Grefenstette, J. J. (1990). GENESIS Version 5. `ftp.aic.nrl.navy.mil:/pub/galist/src`.

Hall, Z. W. (1992). *An Introduction to Molecular Neurobiology*. Sinauer Associates, Sunderland, MA, 1st edition.

Hines, M. L. and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Computation*, 9(6):1179–1209.

Koch, C. and Segev, I. (1998). *Methods in Neuronal Modeling: From Ions to Networks*. MIT Press, 2nd edition.

Lindemayer, A. (1968). Mathematical models for cellular interactions in development. *Journal of Theoretical Biology*, 18:280–299.

Mainen, Z. F. and Sejnowski, T. J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Science*, 382:363–366. `www.cnl.salk.edu/~zach/patdemo.html`.

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

NASA (2000). NASA, Jet Propulsion Laboratory, Revolutionary Computing Technologies: Biological Computing. `cism.jpl.nasa.gov/program/RCT/BioCompUD.html`.

RIKEN (2000). Riken Brain Science Institute. `www.riken.go.jp/`.

Rust, A. G. (1998). *Developmental Self-Organisation in Artificial Neural Networks*. PhD thesis, Department of Computer Science, University of Hertfordshire.

Rust, A. G. and Adams, R. (1999). Developmental evolution of dendritic morphology in a multi-compartmental neuron model. In *Proc. of the 9th Int. Conf. on Artificial Neural Networks (ICANN'99)*, vol. 1, pages 383–388, IEE.

Rust, A. G., Adams, R., George, S., and Bolouri, H. (1998). Developmental evolution of an edge detecting retina. In *Proc. of the 8th Int. Conf. on Artificial Neural Networks (ICANN'98)*, pages 561–566, London. Springer-Verlag.