Evolving the Morphology of a Neural Network for Controlling a Foveating Retina — and its Test on a Real Robot

Peter Eggenberger Hotz[1,2], Gabriel Gómez[1] and Rolf Pfeifer[1]

[1]Artificial Intelligence Laboratory
Department of Information Technology, University of Zurich
Winterthurerstrasse 190, CH-8057 Zuerich, Switzerland
eggen@ifi.unizh.ch

[2]Emergent Communication Mechanisms Project
ATR Human Information Science Laboratories
2-2 Hikaridai, Seika-cho, Soraku-gun
Kyoto 619-0288
Japan

Abstract

The standard approach in evolutionary robotics is to evolve neural networks for control by encoding the parameters of the network in the genome. By contrast, we have evolved a neural controller based on biological principles from molecular and developmental biology. The key principles employed in our algorithms model the specific ligand-receptor interactions and gene regulation. These mechanisms were used to control the growth of the axons, the generation of synapses including the synaptic efficiencies (i.e. the synaptic weights in a neural network model). The evolved neural network was then transferred to a real robotic system with results comparable to the ones achieved the simulation. We hypothesize that the incorporation of mechanisms of gene regulation potentially leads to more adaptive neural networks, that can help bridging the "reality gap" between simulation and the real world.

Introduction

In order to create an artificial evolutionary system (AES) able to evolve complex neural networks for real robots, this paper proposes to combine artificial evolutionary techniques with biological principles. Although this approach is more complicated than a direct encoding scheme or predefined developmental rules, there are a number of advantages of an AES with developmental processes (especially continuous ligand-receptor interactions and genetic regulatory networks) that justify their implementation. In Nature the genome encodes a program of instructions for the development of an organism (Gerhart & Kirschner 1997), (Wolpert 1998). What is stored in the DNA is not so much a detailed description of the body and its parts, but how to build it. The advantage of such an approach is to reduce the amount of genetic information, the number of parameters required to specify the organism. A developmental program is in this sense similar to a data-compression algorithm program. Instead of storing information of each individual point of a picture, it stores instructions on how to re-create the picture with a reduction of information of at least an order of magnitude (Barnsley 1993). A reduction of the number of parameters in an AES for the same problem means higher evolvability and faster convergence towards a solution. In the literature other indirect encoding schemes for artificial evolution were proposed (Fleischer & Barr 1992; Belew 1993; Cangelosi, Parisi, & Nolfi 1994; Gruau & Whitley 1993; Kitano 1994; Michel & Bondi 1995; Sims 1995; Dellaert & Beer 1996; Vaario & Shimohara 1997; Kodjabachian & Meyer 1998). In contrast to these approaches in the proposed model the genes controlled by continuous-valued regulatory genes that control the structure of the neural network as well as the mechanisms for changing the synaptic weights. Different developmental mechanisms (cell division, axonal outgrowth, synaptogenesis, learning) were used to evolve a foveating retina. (By foveating we mean "image tracking"). The retina which is moved by two pairs of antagonistic motors has to learn to focus on an incoming stimulus. The simulations show that the AES with the proposed encoding scheme is able to find solutions for a non-trivial problem, the foveating task, that the genetic code is independent of the number of cells (the number of neurons and synapsis are much bigger than the number of genes) and that the different developmental mechanisms work effectively together. Furthermore we were able to transfer the evolved foveating mechanism on a robot arm guided by a CCD-camera.

The Model

Figure 1 gives an overview of the general levels of the AES. The genetic parameters are defined on the first level and they are directly changed by the evolutionary strategy. These parameters do not encode directly the properties of a neural network, but specify the properties of the cells (second level) by defining the characteristics of the genes and their regulators and they affect the response of the genes. The properties of the third (tissue) and possibly higher levels (organs, organisms) are, again, not directly specified, but have to be created by developmental processes. To do so the evolution strategy was used to explore the parameter space for those processes that are able to fulfil a designer specified fitness function. The use of processes to create structures instead of specifying the structure directly has the advantage that

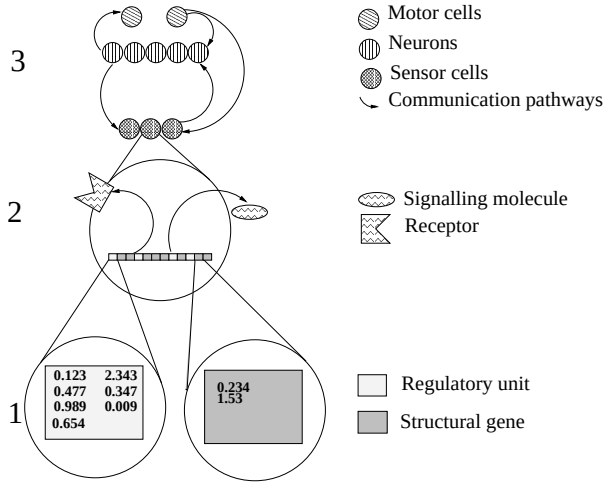

Figure 1: Schematic overview of the artificial evolutionary system (AES). The schema is divided in three levels of increasing complexity. The first level is the parametric level, the second is the cellular level and the third is the tissue level. **Level 1** The parameters of the evolution strategy are grouped in sets of two and seven parameters to define the properties of the regulatory units and the structural genes. **Level 2** The sum of all regulatory units and structural genes is the genome. Although all the cells of a tissue or an organism have the same genome, different sets of genes can be active in different cells and produce different cell types. **Level 3** On the third level the cells have to be able to create structures fulfilling the specified tasks encoded in the fitness function.

the encoding needs less genetic information and therefore decreases the search time in the parameter space. The obvious drawback of this approach is the increased computational load for the development. But as the number of parameters can be decreased drastically there is still an overall advantage, especially for large complex organisms.

Overview of the AES algorithm

Set population size
Set initial gene number
Create n x n x n grid
Repeat
  Initialize first cells
  Diffuse organizing gradients in the grid
  Repeat
    Do for every grid point
      Update cell
        Update gene activities
        Update gene products
        Call developmental processes
      Update environment
        Update diffusing signals
        Update synapses
        Update axons
    End Do
  Until designer set end is reached
  Calculate fitness
  Reproduce and select
    Put fitness in sorting table
    Sort table
    Reproduce the fittest individuals
Until stop criterium is fulfilled

The Structural Genes

Seven parameters (all real valued numbers) encode the properties of each structural gene.

- The first parameter determines what happens if a gene is activated. In the current implementation an active gene either produces a substance such as a signalling molecule or it can activate a predefined function such as cell division. The designer can choose how many of these functions he wants to investigate. Section " Classes of gene products" enumerates which substances and functions were used for the experiments described in this paper.

- The second parameter is used to calculate the probability of an interaction with a partner molecule. This affinity parameter *aff* determines which molecule (signalling molecule or axonal receptor) interacts with which partner (regulatory unit or receptor). To each molecule in the simulation a real valued number and a function is assigned, which calculates an artificial binding affinity between the molecules. This function is implemented as follows:

$$a_{12} = f_{aff}(aff_1, aff_2) = e^{-\alpha(aff_1 - aff_2)^2} \qquad (1)$$

  aff_1, aff_2 are the real valued numbers representing the geometric properties of the substances and are encoded in the parameter set of the evolutionary strategy. α is the affinity parameter with positive values. If α is high, the two substances have to be very similar (i.e.$aff_1 \sim aff_2$) to get a high functional f_{aff} value, if α is low, the substances can be more different to still get high f_{aff} values. Molecules compete for a docking site (cis-regulator, receptor) and their success of binding depends on the affinity between the molecule and the docking site and on the concentrations of the competitors.

- The third parameter determines the sign of the effect e_{ij}, i.e. inhibitory or excitatory.

- The fourth parameter specifies the threshold ϑ. This parameter determines how high the sum of the products of all the affinities a_{ij} between the signalling molecules and the regulators times their concentrations has to be in order to turn on or off the associated structural gene.

- The fifth parameter designates the decay d_i rate for the product.

- The sixth parameter is used to store the affinity parameter α
- The seventh parameter is used as the diffusion parameter D_1

The Regulatory Units

One or several regulatory units control a structural gene. Regulatory units are switches that control the activity of the structural gene. Active regulatory units influence the activity of the structural gene, but only an activated structural gene is able to generate a response such as cell migration or the production of a receptor. Two parameters are assigned to every regulatory unit:

- an affinity aff_{RU}. This parameter has the same use as the affinity parameter in a structural gene. A signalling molecule is defined by the parameters encoded in the structural gene. Both affinity parameters are used to calculate the probability for an interaction between a regulator and a signalling molecule. Both factors are variables of the affinity function $aff_{Tot} = f_{aff}(aff_T, aff_S)$ and its value will influence the probability of a gene's activation.
- threshold. In order to activate a gene, the product of the affinity aff_{Tot} and the concentration of the signalling molecule has to exceed this value.

Gene Regulation

Whether a given gene at position (i, j, k) in a cell on the grid will be activated depends on the affinity and concentration of all the signalling molecules at that position. All these influences are summed up and if this sum exceeds the gene's threshold the gene will be activated or inhibited according to the sign of the effect. All these parameters are varied by the evolutionary strategy and used to explore the interaction space for useful developmental processes able to solve the designer defined tasks. The gene activity of the i-th gene depends on parameters of the structural gene and its regulatory units. The genetic regulatory mechanism is implemented according to the following equations:

$$G_i(c_{sm_0}, \ldots, c_{sm_m}) = \frac{1}{2}(1.0 + \tanh(x)) \tag{2}$$

$$x = \sum_{j=0}(\Theta(a_{ij} c_{sm_j} - \vartheta_j))$$

$$\Theta(x) = \begin{cases} 1.0 & : \text{ if } x > 0 \\ 0.0 & : \text{ otherwise} \end{cases}$$

Where:

- G_i is the activity of the i-th gene
- $\tanh(x)$ is the hyperbolic tangent.
- $a_i j$ affinity to encode the effect between the regulatory unit i and the signalling molecule j (also referred to as transcription factors if they regulate the gene activity).
- c_{SM_j} concentration of the signalling molecules j
- θ_j is a threshold value
- β is a parameter affecting the steepness of Θ

The regulatory units function as reading heads for the concentrations of signalling molecules (inside the cell they are also called transcription factors), which will then determine the state of a gene. The differential equations describing the whole system are:

$$\frac{dg_i(x, y, z, t)}{dt} =$$

$$G_i(c_{sm_0}, .., c_{sm_m}) + D_1 \nabla^2 g_i(x, y, z, t) \tag{3}$$

where:

- $gi(x, y, z, t)$: concentration of substance i at grid point (x, y, z) at time t
- m: number of signalling molecules
- D_1: diffusion constant
- $G_i(c_{SM_0}, .., c_{SM_m})$: see formula (2)
- $D1\ \nabla^2 gi(x, y, z, t)$: diffusion term

Classes of Gene products

The activation of a gene leads to two types of responses: Either a simulated molecule is produced or a function (implemented as a procedure) is executed. The link between the activation of a gene and its response depends on the first of the seven parameters of a structural gene. The following responses were implemented for the experiments in this paper:

1. Production of chemical substances
   (a) a signalling molecule is produced to communicate between the cells.
   (b) a cell adhesion molecule (CAM) is produced to connect the current cell to another one.
   (c) receptors are produced for signalling molecules, axons and synapses.
   (d) synapses bound receptor controlling the synaptic weight
2. The activation of a gene calls a predefined function of the following types:
   (a) cell division

(b) axonal outgrowth
(c) synaptogenesis

The number of different substance classes as well as the number of functions is designer specified and depends on the problem the designer wants to solve. The possible interactions of a chemical substance with a partner molecule depend on the affinity between them (see formula 1) and therefore it is possible that a signalling molecule can interact with different partners.

Inter-cellular communication and cell differentiation

By the gene regulatory mechanism described above the AES is able to simulate cell differentiation based on positional information of diffusible factors or uneven distribution of factors inside the cell. Cell differentiation is the process by which the cells become different because different subset of genes are active in them. Depending on the history of influences a cell has been exposed to, different subset of genes will be active and therefore different cell types will emerge. Note that no genetic information is needed to specify the cell type, because the latter is determined by the inter-cellular communication. (Remember all the cells contain the same genome!). Signalling molecules produced by a cell can diffuse to nearby cells, where they may induce a change of state of some genes, which are influenced differently by the different concentrations of the signalling molecules. Note that this communication mechanism depends on the concentration, the affinity of the signalling molecules as well as the regulatory units of the genes. The reading mechanisms (the regulatory units) vary in their affinity, which explains the different effects on the different cells. The possibility of a cell group to produce a signalling molecule to induce another cell type is an important exploratory mechanism of the AES. It allows exploring new possibilities of communication pathways or neural connections (see Figure 2).

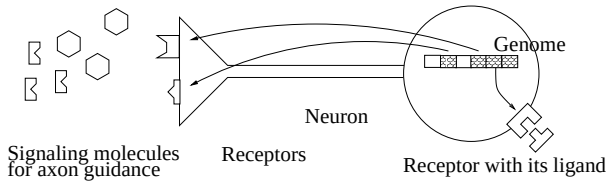

Figure 2: Gene activation. Every cell contains a genome with regulatory units and structural genes. The interactions between regulatory units and signalling molecules determine when a gene becomes active; an active structural gene will then produce a product such as a receptor or will call a function such as axonal outgrowth or cell division.

Neurons

In this section mechanisms pertinent to neural growth and change of synaptic weights are described.

Axonal Outgrowth The interactions between the axonal receptors and the signalling molecules and their respective concentrations determine the direction of the next move an axon will undertake. The effect of the interactions between receptors and signalling molecules at a position (x, y, z) is calculated as follows:

$$I_i = \frac{1}{1 + e^{-(\sum_{i=1}^{n} \omega_{i,j} r_i s_j - \theta_i)}} \tag{4}$$

$I_i = I_i(w, r, s, x, y, z, t)$ is the influence of signalling molecules on the i-th receptor at location (x, y, z) on the grid at time t. The influence is a function of the ligand-receptor interactions and their respective concentrations. r_i is the concentration of the i-th receptor at position (x, y, z) and time t, s_j the concentration of the j-th signalling molecule at position (x, y, z) and time t, $\omega_{i,j}$ the affinity between signalling molecule i and the receptor j and θ_i is a threshold. This formula is calculated in a user-specified range around the current location of the axons, and the next position of the axon on the grid will be in the direction of the biggest influence. The axons are able to read gradients of chemicals and avoid or approach them by reading the concentration differences locally. If an axons finds a cell with matching receptors a link is established.

Neuronal Activity Standard artificial neurons were used in the model and their activities were calculated according to the following formula:

$$S_i = \sigma(\sum_{i=1}^{n} \omega_{i,j} S_j - \Theta) + \text{noise} \tag{5}$$

S_i activity of the i-th neuron, $\omega_{i,j}$ synaptic weight between the i-th and the j-th neuron, Θ is the threshold and σ is the sigmoidal function $\sigma(x) = \frac{1}{1+e^{-x}}$.
The noise term is a randomly added value to allow for spontaneous neural activity.

Control of the Synaptic Weights The change of the synaptic weight depends on specific molecular interactions. The synapses can express receptors whose activity is proportional to the synaptic weight change. One or more receptors can control the weight change and therefore the AES is able to evolve sophisticated learning rules without pre-defining rules how to apply them.

$$\Delta\omega_{i,j} = (\eta \text{rec}_{N_i} \text{rec}_{\text{Syn}_i} \text{rec}_{A_j} - d)\omega_{i,j} \tag{6}$$

Where $\omega_{i,j}$ is the synaptic weight, η the rate of weight change, rec_{N_i} the receptor activity specific for a signalling molecule N_i, $\text{rec}_{\text{Syn}_i}$ the receptor activity for the synapse i, rec_{A_j} represents the activity of neuron j and d the decay rate of the synaptic weight.