Advances In Artificial Life: Proceedings of the 8th European
Conference on Artificial Life, Canterbury, UK, September, 2005

A Coarse-Coding Framework for a Gene-Regulatory-Based Artificial Neural Tissue

J Thangavelautham and G M T D'Eleuterio

Institute for Aerospace Studies, University of Toronto
Toronto, Ontario, Canada, M3H 5T6
thangav@ecf.utoronto.ca, gabriele.deleuterio@utoronto.ca

Abstract. A developmental Artificial Neural Tissue (ANT) architecture inspired by the mammalian visual cortex is presented. It is shown that with the effective use of gene regulation that large phenotypes in the form of Artificial Neural Tissues do not necessarily pose an impediment to evolution. ANT includes a Gene Regulatory Network that controls cell growth/death and activation/inhibition of the tissue based on a coarse-coding framework. This scalable architecture can facilitate emergent (self-organized) task decomposition and require limited task specific information compared with fixed topologies. Only a global fitness function (without biasing a particular task decomposition strategy) is specified and self-organized task decomposition is achieved through a process of gene regulation, competitive coevolution, cooperation and specialization.

1 Introduction

Evolving open-ended variable-length neural systems with large phenotypes remains a significant challenge in the field of Alife [20, 8]. One of the problems encountered with large phenotypes is the bootstrap problem [18]. The bootstrap problem occurs when the EAs are unable to pick out incrementally better solutions for crossover and mutation resulting in premature stagnation of the evolutionary run. The answer to the evolution of controllers for complex problems has often been to introduce more supervision *ad hoc*, where the experimenter decomposes a complex task into a set of simpler tasks based on domain knowledge of the task at hand. In biological systems, such intervention (more supervision) does not always exist, yet these systems can adapt and thrive with relative ease.

Other techniques involve starting with a single cell or a small phenotype and allowing for the system to grow in size and complexity until the system can find a satisfactory solution to a given task [14]. However in biological systems, often there exists a brain that may already have the neural capacity (brain size) to adapt easily to a new task/scenario. In such circumstances, starting over with a minimalist system may be a much slower process owing to the over-reliance of topological growth directed by evolution.

Taking inspiration from the mammalian visual cortex, we have developed an evolvable Artificial Neural Tissue (ANT) model. The genotype for the evolvable

ANT defines a developmental program that constructs a neural tissue (phenotype) and associated gene-regulatory functionality. The variable-length tissue architecture can be characterized as a lattice of neurons arranged in a three-dimensional (3-D) structure. A Gene Regulatory Network (GRN) controls cell division and cell death in the tissue, activates/inhibits portions of the tissue based on external sensory input using a coarse-coding framework and express/repress other characteristics based on gene-protein interactions.

We empirically compare the training performance (using evolutionary algorithms) of various controllers for the multirobot tiling pattern formation task (similar to a segment of the termite nest building task [4]),variant of the phototactic task (with obstacles and robot equipped with a gripper) and a relatively difficult sign-following task that requires use of memory. Each of these tasks requires the evolution of emergent (self-organized) task decomposition strategies to complete the task successfully. Only a global fitness function (without biasing a particular task decomposition strategy) is used; the intention is for the controllers to evolve innovative techniques in decomposing the global task into a set of 'local' subtasks.

2 Related Work

Traditional machine learning methods for task decomposition involve the supervisor decomposing the task and training separate 'expert networks' on the subtasks [12]. Arbitration among the expert networks is performed using a cooperative (Product of Experts model) [10] or competitive mixture of experts model [12, 17]. The gating function and the expert networks are trained separately and the network topology is predetermined by the experimenter.

Our previous work took this approach to the next step, where decision neurons (acting like gating functions) and expert networks were evolved together (Binary Relative Lookup architecture) using a global fitness function [21]. We found larger BRL architectures (with more expert networks) tend to evolve faster than comparable smaller ones for the tiling pattern formation task. The decision neurons learned to limit the number of expert networks used thus preventing problems in over segmentation (over-fitting) to many expert networks.

NEAT (NeuroEvolution of Augmenting Topologies) showed the potential advantage of evolving both the neuron weights and topology concurrently [13]. It is argued that growing the network incrementally ('complexification') serves to minimize the dimensional search space and thus improve evolutionary performance. ANT is even more flexible and can be initialized with a large number of neurons since the GRNs can effectively suppresses unnecessary/noisy neurons while activating neurons specialized for specific input signals.

Another approach to evolving solutions to complex tasks involves use of encoding schemes that effectively reduces the search space. This includes a multicellular developmental systems by Eggenberger [5] and the Morphogenetic System (MS) originally used on POEtic [19]. Eggenberger's earlier model demonstrates how cell differentiation, cell division and cell death can be controlled by gene reg-

ulatory functionality and constructs a 3-D organism. In these two systems, the GRNs merely act on the developmental process in constructing the phenotype.

A more refined model by Gomez and Eggenberger [6] uses 'ligand-receptor interactions' allowing for one neuron to recognize/attach to a partner neuron and allow for emergence of Hebbian-type learning without specification of learning rules for a forveating artificial retina system. Astor and Adami's [2] tissue architecture consists of cells on a 2D tissue that perform logical functions. Cell replication and connections are formed through a gene regulated developmental and learning system using a Genetic Programming type command set.

In our ANT architecture, gene regulation occurs during the developmental process in addition to when the tissue interacts with the environment. The decision protein act much like gating neurons while helping to reduce the effects of spatial crosstalk [12] and perform sensory preprocessing enabling selection of 'specialized' networks of neurons depending on the sensory input.

Another class of indirect developmental encoding schemes such as Artificial Embrogeny systems [14] produce phenotypes through recursive rewriting of the genotype code. These systems use an artificial chemistry as a grammar or model cellular metabolism and replication. Other recursive rewriting schemes include Cellular Encoding [7] and L-Systems (see [15, 11]).

It has been argued that indirect developmental encoding schemes may effectively decrease the search space (by exploiting regularities and allowing for peleiotropy) but at the price of introducing a deceptive fitness landscape [20]. It has also been found the overhead required for indirect encoding schemes appear to result in poor performance for smaller search spaces [20]. This presents a problem for task decomposition, where the control scheme needs to work well for subtasks with small and large search spaces.

3 Methodology

The ANT architecture presented in this paper consists of a developmental program that constructs a neural tissue and associated gene-regulatory functionality. The gene-regulatory network consists of parameters that control growth and activates/inhibits parts of the genotype. All the parameters defined within the ANT architecture are evolved. This includes parameters characterizing the decision proteins, growth parameters, cell contents and tissue topology. Neural networks consisting of cells within the tissue are dynamically formed through a sequence of activation/inhibition based on the coarse-coding framework.

The artificial tissue consists of a culture of cells activated and inhibited by a gene-regulatory network. The cells exist in a three-dimensional matrix with each cell occupying a cube (Fig. 2). Each cell contains genes specifying weights, thresholds/biases, choice of activation function (modular neuron model [21]) and probability ratios for instructing cell division. Cell division requires a parent cell (selected with highest replication probability using GRNs). The new cell can be located in one of 6 neighbouring locations (top, bottom, north,south,east,west) sharing a common side with the parent and is not occupied by another cell.

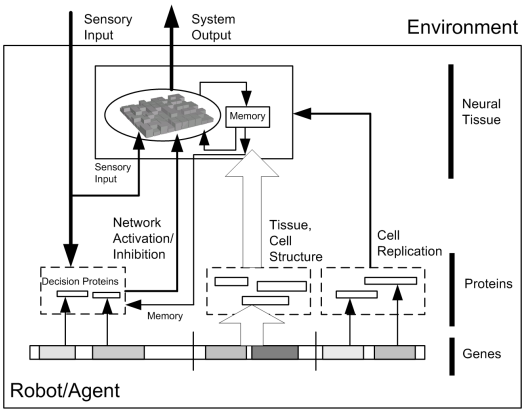

Fig. 1. Schematic of the ANT architecture showing gene-protein-tissue interaction

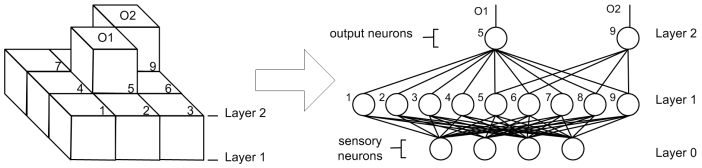

Fig. 2. Diagram of tissue morphology and equivalent networks where O1 and O2 are output signals.

The base layer of cells within the tissue are fully connected to all the sensory neurons or memory neurons. While the top layer of cells triggers a set of predefined basis behaviours such as to perform motor control or pass data to memory neurons. Connection between cells from one layer to another is local, with each cell from layer z connected to a maximum of $m = 9$ cells from layer $z-1$ (spatially localized). $p_{x,y,z} = (\sum_{i=x-1}^{x+1}(\sum_{j=y-1}^{y+1} w_{i,j,z-1}s_{i,j,z-1}))/(\sum_{i=x-1}^{x+1}\sum_{j=y-1}^{y+1} s_{i,j,z-1})$ and $s_{x,y,z} = [\phi_n(p,t_1,t_2)]_{x,y,z}$ where $w_{x,y,z}$ is a neuron weight and $s_{x,y,z}$ is the current state of a neuron. The modular neuron model used consists of two threshold parameters t_1 and t_2, where each neuron outputs one of two states $s = (s_1, s_2)$. A choice of four threshold activation functions for ϕ_n is given below:

$$\begin{aligned}
&\phi_1: \; s_{out} = \begin{cases} s_1, & \text{if } p \le t_1 \\ s_2, & \text{if } p > t_1 \end{cases} \qquad \phi_3: \; s_{out} = \begin{cases} s_1, & \text{if } t_2 < p < t_1 \\ s_2, & \text{otherwise} \end{cases} \\
&\phi_2: \; s_{out} = \begin{cases} s_1, & \text{if } p \ge t_2 \\ s_2, & \text{if } p < t_2 \end{cases} \qquad \phi_4: \; s_{out} = \begin{cases} s_1, & \text{if } p > (1-p) \\ \text{rand}(s_1,s_2), & \text{if } p = (1-p) \\ s_2, & \text{if } p < (1-p) \end{cases}
\end{aligned} \tag{1}$$

enabling a single neuron to simulate AND, OR, NOT and XOR functions.

3.1 Decision Proteins and Coarse Coding

Albus argued that the mammalian brain uses tile (coarse) coding to represent multidimensional space [1, 9]. Hinton [9] and Ballard [3] point to the importance

of modularity in the coarse-coding framework as an increase in the number of neurons in a fixed volume limits the number of dimensions represented. We show in our model, the potential advantages of coarse-coding as a means of arbitration between modular networks.

The activation and inhibition of genes occur through a coarse-coding framework. Decision proteins (modeled as single neurons with ability to choose between threshold activation functions) get activated and inhibited due to sensory input. The decision proteins act by diffusing through a coarse column (receptive field) as shown in Fig. 3. The receptive field parameters specifying position and dimensions ($D_i[x, y, x_{length}, y_{length}]$) for each decision protein D_i is also evolved.

Once a decision protein is activated, the activation concentration $C_i[c_{active}]$ of each neuron cell i is incremented by k (a constant). With multiple decision proteins acting together, a consensus is reached when the activation concentration (summed over multiple activated proteins) is highest for a particular output neuron (cell at the top layer of a column). A network is dynamically formed from all the neurons connected to the selected output neuron (in a scheme shown in Fig. 3) with activation concentration ($c_{active} > 0$). This characteristic is inspired from the columnar pooling of neurons within the mammalian visual cortex. Similarly, if there are multiple output neurons with the same activation concentration, selection occurs among the output neurons according to a uniform distribution.

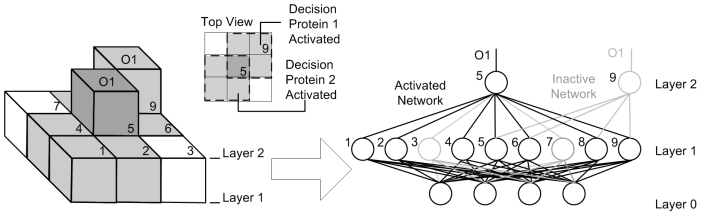

Fig. 3. Diagram showing network with neuron 5 (layer 2) with highest activation concentration (due to coarse activation of decision protein 1 and 2) being selected.

3.2 Introns

The accumulation of nonsense genes or introns has been described to be a major problem in the field of Genetic Programming. However, introns allow for genetic neutrality, an important facet of evolution. With current GP approaches, intron accumulation results in increased computational inefficiency over time. Current techniques in controlling introns has been to prune the genotype regularly or explicitly impose a size limit on the genotype using the fitness function. Both strategies either lack biological plausibility or translate into more supervision.

It is noted that with GP approaches, program size increases according to the square power law in generations [16]. For ANT, the program size growth is controllable and increases linearly due to the growth model used. Our genotype-to-phenotype mapping scheme avoids intron accumulation due to the crossover operator. In GP, a node is chosen from each genome and genetic code is exchanged about the node. In our methodology , only a 'compatible set' of genes

get interchanged during crossover. Each cell has a unique position parameter $[x, y, z]$ relative to rest of the cells within the tissue. A crossover is performed by drawing a plane (with normal vector parallel to the x or y-axis) separating the tissue and exchanging 'compatible' cells (Fig. 4). Thus genes for cell C_1 from Tissue A and C_2 from Tissue B could be interchanged iff $C_{A,1}[x, y, z] = C_{B,2}[x, y, z]$.

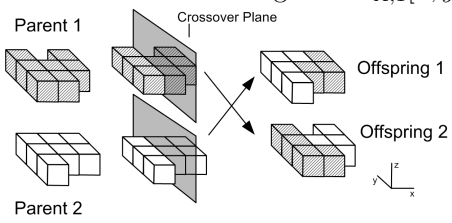

Fig. 4. A crossover operation between two ANT parents. 'Compatible' cells are interchanged as shown resulting in two offspring.

4 Example Tasks

The evolutionary performance of our ANT architecture is compared with fixed network topologies (direct-encoding schemes) for three different robotic tasks. All three robotic tasks were chosen because it could be argued that self-organized task decomposition strategies may be necessary to accomplish the tasks given a 'global' fitness function. In addition, these tasks are inspired by some remarkable behaviours evident in the insect world. The tasks include a multirobot tiling pattern formation task [21] that involves redistributing objects (blocks) randomly placed in a two-dimensional world into a desired tiling structure (see Fig. 7). The robots need to come to a consensus and form one 'perfect' tiling pattern. This task is similar to a segment of the termite nest construction task that involves redistributing pheromone filled pellets on the nest floor [4].

The tiling formation task may be decomposed into a set of subtasks such as foraging for blocks, redistributing block piles, arranging blocks in the desired tiling structure locally, merging local lattice structures, reaching a collective consensus and finding/correcting mistakes. In the phototactic task, the robot must reach a goal location (light source), but it also needs to negotiate obstacles.

In the sign-following task, the robot needs to evolve the ability to decipher the signs relative to the robot's current frame of reference,to remember the current sign while looking for the next one, and negotiate obstacles. Each sign is a waypoint that gives direction to the next waypoint leading ultimately to a goal location. Of the three tasks, this task is the most complex and requires the use of memory. This task is inspired by honey bees ability to waggle (communicate with other bees) and describe directions and waypoints to a food source.

The signs are posted on a fixed frame of reference and the robot based on its current heading needs to interpret the signs accordingly. For the phototactic task and the sign-following task, the fitness function is simply the number of times the robot reaches and stays at the goal location and for the tiling formation task the fitness is the Shannon's entropy over all the tiles [21].

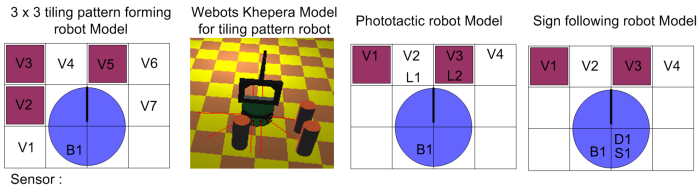

Fig. 5. Schematic of 2D robot model. B1 detects whether robot is carrying a block, D1 is a compass sensor and S1 reads the colour of the sign off the floor .

The robots are modeled as Kheperas equipped with a gripper turret. We have developed a fast two-dimensional grid world simulator for our robotic experiments and we have verified the performance of the evolved controller for the phototactic and tiling pattern formation task using Cyberbotic's Webots (Fig. 5). The basis behaviours chosen for all three tasks are shown in the table below:

Task	3x3 tiling pattern formation	Phototactic	Sign Following
Basis Behaviours	Move Pickup/Putdown NOP	Pickup/Putdown Move Forward Turn Right Turn Left NOP	Pickup/Putdown Move Forward Turn Right Turn Left NOP
Memory Neurons*		2	4
Monolithic Network Topology	4 hidden neurons, 1 output neuron	12 hidden neurons, 4 output neurons (without memory) 8 output neurons (with memory)	36 hidden neurons, 12 output neurons

Table 1. Basis behaviours for the three robotic tasks and monolithic network topology shown. A combination of these behaviours can be activated at each timestep.

4.1 Experiments

The evolutionary performance of various control system architectures is compared for the three tasks (see Fig. 6). The fixed network architectures (Table 1) were determined based on the number of neurons required for triggering each basis behaviour and the size of the networks were limited to avoid the 'bootstrap problem'. BRL2 consist of 2 monolithic networks (Table 1) arbitrated by a decision neuron. ANT is initialized with individuals between 40 to 400 cells (uniform distribution) for the tiling formation task and 20 to 100 cells for the phototactic and sign-following task. This procedure is intended to determine wether there is evolutionary preference for smaller phenotypes over larger ones.

The EA population size for all three experiments is $P = 100$, crossover probability $p_c = 0.7$, mutation probability $p_m = 0.025$ and tournament size of $0.06P$ (for tournament selection). For the phototactic task and the sign-following task the success rate is averaged over 100 different initial conditions (consisting of 20×20 world, 80 blocks) and for the tiling formation task, the fitness is averaged over 15 initial conditions. For the sign-following task 'mines' (undetectable by robot) are randomly laid throughout the gridworld except along the pathway.

5 Results and Discussion

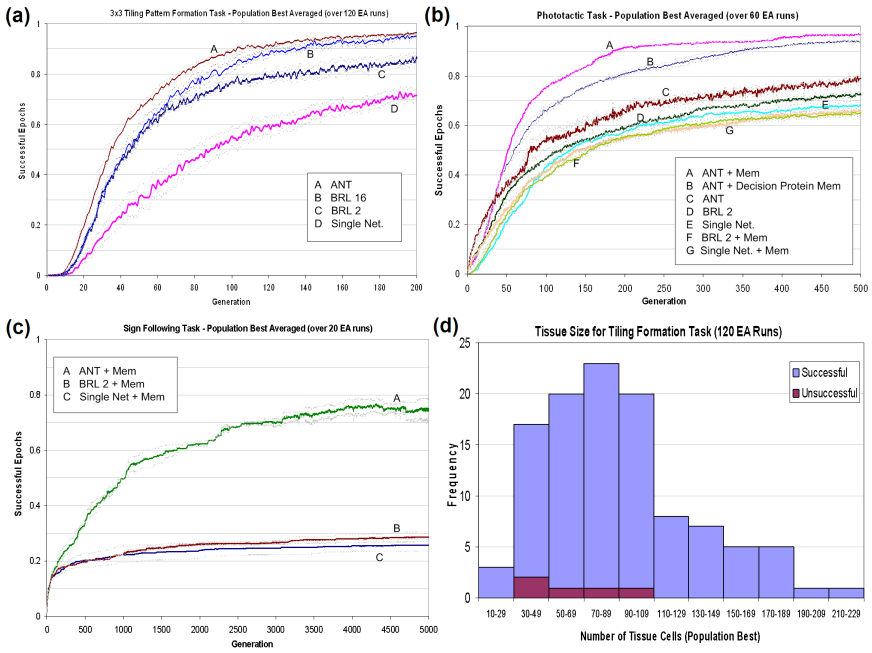

Fig. 6. Evolutionary performance comparison for the tiling-formation (a), phototactic (b) and sign-following tasks(c). Also shown, a histogram (d) of number of cells within the tissue (population best, after 200 generations), for the tiling formation task.

The evolutionary performance (number of successful epochs) of the ANT architecture for three different robotic tasks is better than smaller fixed network topologies (Fig. 6). In addition, ANT with its ability to grow in complexity (depending on task) and exploit modularity managed to find solutions to the sign-following task (memory dependent) where fixed topologies appear to fail.

Analyzing the 3-D morphology, active segments (consisting of specialized networks) are distributed sparsely throughout the tissue. These networks do not appear to decompose the task according to 'recognizable' distilled behaviours but as a set of emergent proximal behaviours (proximity in sensor space) [18].

Large tissue structures do not seem detrimental to the evolutionary performance because the GRN quickly evolves the ability to suppress unnecessary/noisy neurons within the tissue. There appears to be a steady reduction in the number of active neurons with a convergence to a solution (Fig. 7b). During the early phase of evolution, there is greater network activity enabling sampling of more neurons. Convergence appears to result in the inhibition of cells that add noise to the system (reduction in spatial crosstalk [12]).

There appears to be a noticeable improvement in evolutionary performance of large tissue structures over smaller ones. A histogram of population best during successful runs (with a success rate of > 0.9) compared to population best during unsuccessful runs (success rate < 0.1) appear to confirm these observations (Fig.

6d). The ANT architecture is more effective at being able to decompose a task and handle the subtasks using simpler specialized networks consisting of few neurons than the fixed architectures. This results in fewer weights having to be tuned and thus a smaller search space.

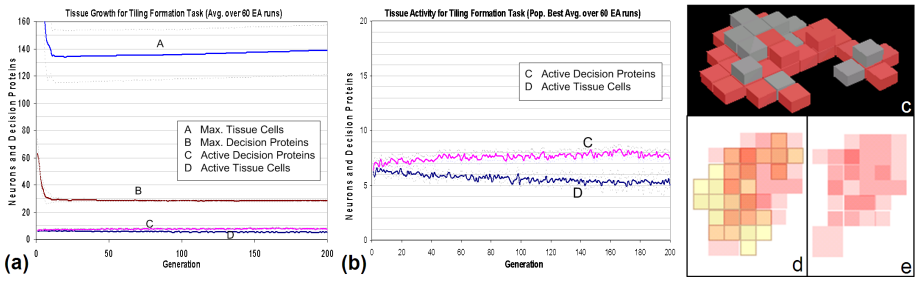

Fig. 7. Tissue growth characteristics (a) and max number of active cells/decision proteins (b). (c) 3D morphology of a tissue (solution for tiling formation task). Top view (d) showing all the decision proteins/cells and top view (e) of decision proteins.

The most promising result from our experiments is the successful exploitation of global memory by the ANT architecture. While a fully connected feed forward network and the BRL architecture were unable to exploit this functionality resulting in lower performance (due to a larger search space). This observation corroborates Nolfi's experiments for the garbage collection task [18].

The ANT architecture consists of cells interconnected locally but with access to global memory, specialized networks separated by distance have a means of communication, resulting in output signals that are maybe less contradictory. Access to global memory allows for some emergent (self-organization) properties to appear with one segment of the tissue able to 'veto' or override another segment of the tissue independently. This is evident from significant performance improvement for the phototactic task when only the decision proteins were able to read the memory neurons (Fig. 6b).

It is also observed that the number of active decision proteins steadily increases even after convergence to a solution thus creating a redundant protective mechanism against adverse mutation (Fig. 7b). Redundancy is an important characteristic of biological systems that enable a system as whole to be robust despite the fragility of its parts. From an evolutionary standpoint, the artificial tissues evolve not only to solve the intended task, but also steadily develop ways to ensure its survival and successful transfer of its genes.

6 Conclusion and Future Work

A developmental Artificial Neural Tissue architecture, with the ability to perform emergent task decomposition is presented in this paper. The evolutionary performance of the ANT architecture is better than smaller fixed network topologies for three different robotic tasks. Our experiments indicate improved evolutionary performance of ANT with access to memory neurons. It is hypothesized that memory neurons allow for communication among the specialized networks,

enabling one network to override another. We are currently planning to port ANT onto hardware and hope to compare our architecture with other variable length topologies for tasks such as soccer, mining and surface exploration.

Acknowledgements. The authors would like to thank NSERC and the Ontario Government for supporting this research and to the reviewers for their helpful comments.

References

1. Albus, J. S. A theory of cerebellar function. Mathematical Biosciences, (1971) 25–61.
2. Astor JC, Adami C. 2000. A developmental model for the evolution of artificial neural networks. ALife 6(3):189-218
3. Ballard, D.H., Cortical Connections and parallel processing. The Behavioural and Brain Sciences, Vol. 9, (1986) 279–284
4. Deneubourg J-L. Application de l'ordre par fluctuations 'a la description de certaines 'etapes de la construction du nid chez les termites. Insectes Sociaux, (1977) 117–130
5. Eggenberger, P., et al., Evolving the morphology of a neural net for controlling a foveating retina, ALife 8 (2002) 243-251
6. Eggenberger, P., Evolving Morphologies of simulated 3d organism based on Differential Gene Expression, Proc. 4th European Conference on ALife (1997)
7. Gruau, F. Automatic definition of modular neural networks. Adaptive Behaviours 3 (1994) 151–183
8. Gustafson S., et al.,Problem Difficulty and Code Growth in Genetic Programming Genetic Programming and Evolvable Machines 5, (2004) 271–290
9. Hinton, G., Shape Representation in Parallel Systems, Proc. of 7th Int. Joint Conf. on Artificial Intelligence (1981) 1088–1096
10. Hinton, G.,Product of Experts, Proc. of the 9th Int. Conf. on Artificial Neural Nets (1999) 1:1-6
11. Hornby, G., Pollack, J. Creating High-Level Components with a Generative Representation for Body-Brain Evolution. ALife 8, 2002
12. Jacobs R., Jordan M., Barto A., Task decomposition through competition in a modular connectionist architecture. Cognitive Science, No. 15, (1991) 219-250
13. Stanley K.,Miikkulainen R.,Continual Coevolution through Complexification in Proc. of the Genetic and Evolutionary Computation Conf. 2002, (2002)
14. Stanley K.,Miikkulainen R., A taxanomy for aritificial embrogeny., Artificial Life 9, (2003) 2:93-130
15. Sims K., Evolving 3D Morphology and Behavior by Competition, Proc. of Artificial Life IV, (1994) MIT Press, 28–39.
16. Langdon, W., Quadratic bloat in genetic programming, in Proc. of Genetic and Evolutionary Comp. Conf. (2000)
17. Liu Y.,Yao X., Higuchi T., Evolutionary Ensembles with Negative Correlation Learning. IEEE Transactions on Evolutionary Computation (2000) 4:380-387
18. Nolfi, S., Floreano D.: Evolutionary Robotics : The Biology, Intelligence, and Technology of Self-Organizing Machines, MIT Press (2000) 13-15
19. Roggen D., Floreano, D., Mattiussi, C., A Morphogenetic Evolutionary System: Phylogenesis of the POEtic Tissue, Int. Conf on Evolvable Systems (2003) 153–164
20. Roggen D., Federici D., Multi-cellular Development: Is There Scalability and Robustnes to Gain?, In Proc. of Parallel Problem Solving from Nature (2004) 391–400
21. Thangavelautham, J.,D'Eleuterio, G.M.T, A Neuroevolutionary Approach to Emergent Task Decomposition, Proc. of 8th Parallel Problem Solving from Nature (2004) 991-1000