Developmental Neural Heterogeneity through Coarse-Coding Regulation

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 coarse-coding regulatory model that facilitates neural heterogeneity through a morphogenetic process is presented. The model demonstrates cellular and tissue extensibility through ontogeny, resulting in the emergence of neural heterogeneity, use of gated memory and multistate functionality in a Artificial Neural Tissue framework. In each neuron, multiple networks of proteins compete and cooperate for representation through a coarse-coding regulatory scheme. Intracellular competition and cooperation is found to better facilitate evolutionary adaptability and result in simpler solutions than does the use of homogeneous binary neurons. The emergent use of gated memory functions within this cell model is found to be more effective than recurrent architectures for memory-dependent variants of the unlabeled sign-following robotic task.

1 Introduction

One of the big challenges in Alife is to design open-ended artificial multicellular developmental systems that can grow in complexity to solve extensible control tasks by performing task-decomposition with little or no explicit supervision. Fundamental to understanding and re-engineering multicellular biological systems is to determine how functionality is distributed within these systems and how specialization takes shape.

It has been theorized that exploratory selection/regulation mechanisms, the process by which selection of parallelized selection of functional outcomes facilitates *evolutionary adaptability*, that is, the ability for genes to be heritable and selectable phenotypes, less susceptible to lethal mutations and produce novel traits with fewer mutations [12]. Extensive evidence of exploratory selection processes has been found in the immune system and this had spurred interest into how these processes might be at work within the brain [3, 12]. Furthermore, biological evidence hardly points to the notion of a 'typical' homogeneous feedforward binary neurons of the McCulloch-Pitts type. Neurons are, in fact, complex heterogeneous multistate analog systems with memory.

In this paper, we demonstrate the advantages of these exploratory selection/regulation mechanisms based on a coarse-coding scheme, inspired by Albus [1], and show evidence of emergent task decomposition and specialization

occurring at the cellular and gene/protein level within an Artificial Neural Tissue (ANT) framework [16]. The model exhibits cellular and tissue extensibility through ontogeny, resulting in the emergence of neural heterogeneity, use of memory and multistate functionality within the ANT framework. Within each neuron, multiple networks of protein compete and cooperate for representation through a coarse-coding framework for binding sites. We choose to use coarse coding as it is a moderately distributed coding scheme that allows for pooling and redundancy thus helping to render the system robust in the face of noisy sensor data.

This model, with no explicit supervision and limited task-specific assumptions, produces solutions to a variant of the sign-following task found to be unsolvable with fixed-topology homogeneous artificial neural networks. Fixed-topology networks that lack regulatory functionality perform poorly in complex tasks with limited supervision owing to the bootstrapping problem, which causes premature stagnation of an evolutionary run [13].

2 Background

Artificial developmental systems mimic ontogenic processes from biology and have been successfully used, with variable-length genomes, to 'grow' topologies and heterogeneous functionality without explicit intervention. ANT is a morphogenetic system with a directly encoded genome and uses gene regulatory systems (GRNs) for development. Artificial embryogenic systems (L-systems [11] and cellular-encoding systems [8]) use indirect encoding schemes that involve recursive rewriting of the genotype to produce a phenotype. However, it has been argued that indirect encoding schemes introduce a deceptive fitness landscape and result in poor performance for smaller search spaces owing to overhead [15].

Examples of artificial morphogenetic systems include the work by Eggenberger [4] and by Gomez and Eggenberger [5], the latter using 'ligand-receptor interactions' to perform cell aggregation. A morphogenetic system was also used on POEtic by Roggen *et al.* [14]. Developmental tissue models such as Norgev by Astor and Adami [2] are also morphogenetic and facilitate cellular heterogeneity. Cell replication and synaptic connections are formed through a GRN based developmental and learning system using a genetic-programming-type command set. However, in our ANT model, regulation continues after development at the gene/protein and cellular levels. Neuroregulatory functionality is performed through coordinated release of diffusive neurochemicals resulting in superpositioning of chemical concentration fields (Fig. 1). Other models such as GasNet allow for volume signaling between neurons using neurochemicals but lack explicit regulatory functionality [10].

In the coarse-coding cell model presented here, a selection/regulation process is also at work within protein networks resident in each neuron, thus facilitating heterogeneity and open-ended growth in complexity of cells. The use of Cellular heterogeneity may be more biologically plausible but more complex, multistate

cells do not necessarily present advantages over simpler binary-state cells. Yet, the use of multistate feedforward and memory neurons may be beneficial.

The need for specialized memory neurons arises from the 'error decay' problem evident with standard recurrent connections for learning processes. A stored signal remains unprotected from spurious inputs decays or grows (without bound) making it difficult to recall a signal after many timesteps [6]. Long Short-Term Memory (LSTM) [9] overcomes this limitation but it is a predefined architecture consisting of a storage neuron, a reset gate, an input gate (to protect memory from spurious inputs) and an output gate. However, for the T-maze task (a simpler variant of the sign-following task), it was found that recurrent networks trained using Enforced Subpopulations (ESP) and Hierarchical ESP (H-ESP) outperformed a LSTM architecture [7]. LSTM also lacks biological plausibility and similar memory functions can be obtained without the use of predefined cell blocks using the coarse-coding cell model.

3 Artificial Neural Tissue and Coarse-Coding Cell Model

The ANT architecture consists of a developmental program, encoded in the genome that constructs a three-dimensional neural tissue and associated regulatory functionality (see [17] for details). The tissue consists of two types of neural units, decision neurons and motor-control neurons. Regulation is performed by decision neurons, which dynamically excite or inhibit motor-control neurons within the tissue based on a coarse-coding framework (Fig. 1).

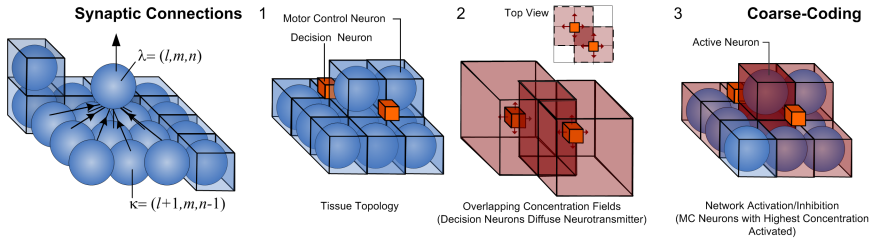

Fig. 1. Synaptic connections between motor-control (MC) neurons and operation of neurotransmitter field.

Our neuron model permits a number n_c of messenger-channel protein networks (Fig. 2a). Each protein network receives the same inputs $\mathbf{x}$, an $n \times 1$ real-valued column, which represents either sensory data or inputs from other neurons. These inputs are fed through m_j 'ion channels' that transform 'electrical signals' into various types of ion, the concentration of which are collectively denoted $\mathbf{y}_j$, an $m_j \times 1$ real-valued column. The concentrations are given by

$$\mathbf{y}_j = \mathbf{W}_j \mathbf{x}$$

where $\mathbf{W}_j$ is an $m_j \times n$ real-valued weight matrix associated with the jth protein network.

Each protein network produces an 'activation protein,' whose concentration c_j is determined by a linear combination of p_j basis functions $\psi_{ik}, i = 1 \cdots m_j, k = 1 \cdots p_j$, dependent on the ion concentrations $\mathbf{y}_j$. The ith basis function, in fact, depends only on the ith ion concentration y_{ij}:

$$\psi_{ik}(y_{ij}) = \begin{cases} 1, & \text{if } \tau_{1,k} \leq y_{ij} \leq \tau_{2,k} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

The boundary parameters, τ_1 and τ_2 for each ψ are evolved. The concentrations c_j that determine the output of the neuron are simply $c_j = \sum_{i=1}^{m_j} \sum_{k=1}^{n_{b,j}} \psi_{ik}(y_{ij})$. (Note that c_j are integers.) This structure is reminiscent of and was inspired by the coarse-coding scheme of Albus [1] and we accordingly refer to it as coarse-coding regulation. The basis functions here square-hat functions in one dimension although the kind of tiled functions in Albus's Cerebellar Model Arithmetic Computer can also be used.

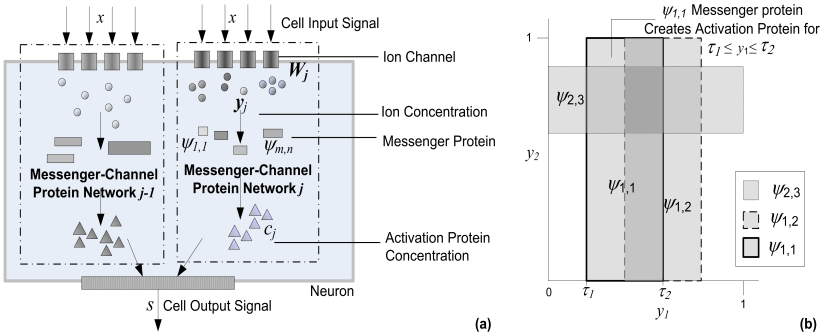

Fig. 2. (a) Schematic of competing messenger-channel protein networks. (b) Coarse-coding interactions between messenger-protein for $m_j = 2$.

Binary-State Neuron. Let us first consider a binary-state neuron, i.e., one where the output $s(t)$, where t represents the discrete time-step, is either 0 or 1. This output is given by

$$s(t) = \begin{cases} a_j, & \text{if } c_j \text{ is a unique maximum} \\ \phi s(t-1), & \text{otherwise} \end{cases} \tag{2}$$

The messenger-channel protein networks compete to determine the neuron's output. If $c_j = \max\{c_1, c_2 \ldots c_{n_c}\}$ and is uniquely determined, i.e., no two networks produce the same maximum concentration, then the output is taken as $a_j \in \{0, 1\}$ (genetically evolved). Otherwise, the output takes the value $\phi s(t-1)$, where $\phi \in \{0, 1\}$. When $\phi = 1$, the output from the previous time-step is maintained. Thus $s(t)$ is intended to model the 'spike' status of the neuron.

Multistate Neuron. Spiking neurons superimpose their spikes on a background signal. We model this aspect of the neuron by allowing for a multiple-state output, $\mathbf{s} = [s_1 \;\; s_2]$ where $s_1 \in \{0, 1\}$ is associated with the spiking signal and s_2 with the background signal. The output s_1 is given again by (2). We offer two

models for the computation of s_2, a feedforward model and a memory model. In the former, s_2 is given by (2) with a_j replaced by $b_j \in \{0, 1/q_b, ..., 1\}$ which is graduated in q_b (an integer greater than one) uniform steps between 0 and 1. In the memory model, $b_j \in \{0, \max\{0, s_2(t-1) - 1/q_b\}, \min\{s_2(t-1) + 1/q_b, 1\}, 1\}$ allowing for storage, reset, gating and increment/decrement functionality. This multistate model of the neuron is an attempt to better incorporate biological observations of neuron action potential through bottom-up modeling of protein interactions.

Evolution and Development Details of the development process for the tissue remains identical to previous versions of ANT and can be found in [16, 17]. Unlike previous versions of ANT that used neurons with a modular activation function using two thresholds [18], the coarse-coding cell model allows for a developmental activation function. Cell and protein genes have a binary 'activation' parameter, used either to express or repress gene contents. (The genome structure is shown in Fig. 3.) Each channel protein references a cell address. Messenger and action proteins in turn reference a channel protein. Since these genes are modular, it is possible for a messenger-channel protein network to be incomplete and thus lacking channel proteins.

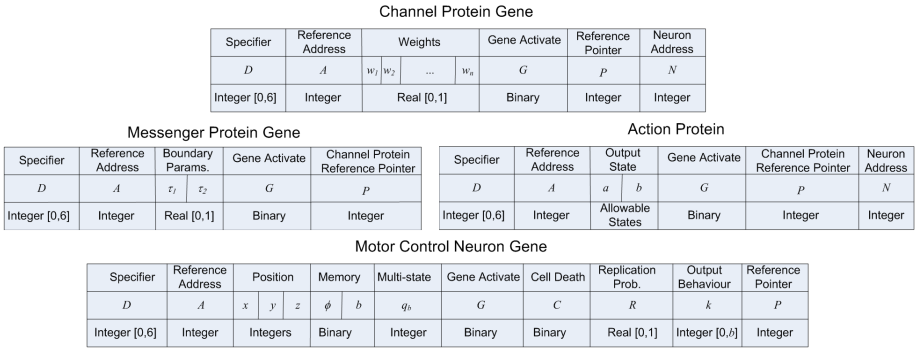

Fig. 3. Genome of messenger-protein network components and a typical motor control neuron.

Mutations in the genome can perturb existing genetic parameters or addition of new (cell, messenger, channel or action) genes caused by random gene transcription errors with a probability of p_{te}. Thus a new cell-protein gene as a result of a transcription error is a copy of an existing cell-protein gene with perturbations starting at point chosen from a uniform distribution along the gene's length and with the gene activation parameter toggled off by default.

4 Sign-Following Task

The effectiveness of the coarse-coding cell model is demonstrated in simulation on two memory-dependent versions of the unlabeled sign-following task. The

workspace is modeled as a two-dimensional grid environment with one holonomic robot (based on a Khepera, equipped with a gripper and camera) occupying four grid squares. For these tasks, the controller must possess a number of capabilities including that to decipher signs relative to the robot's current frame of reference, to remember the current sign while looking for the next one, and to negotiate obstacles (see Fig. 4a). Each sign is color-coded and represents a waypoint (posted in a fixed frame of reference) that gives direction in one of four cardinal points to the next waypoint leading ultimately to the goal location.

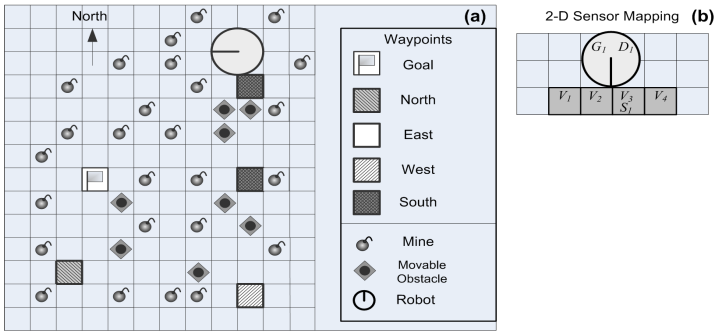

Fig. 4. (a) 2D grid world model for the sign-following tasks. (b) Input sensor mapping.

Mines (undetectable by the robot) are randomly laid throughout the floor except along the pathway. Once a robot encounters a mine, it remains disabled for the remainder of its lifetime. The sensory input map is shown in Table 1 (see also Fig. 4b). The task has to be accomplished using a discrete set of basis behaviors specified in Table 2. These behaviors are activated based on controller output and all occur within a single time-step. The robot is initially positioned next to the first sign, but the initial heading is randomly set to one of the four cardinal directions. Since the robot can only detect signs placed in front, it needs to go into a 'sign searching' mode and perform a sequence of 'turn left' or 'turn right' behaviors to detect the first sign. Once the first sign is detected, the robot then needs to transition to a 'sign following' mode, requiring one bit of memory.

Table 1. Sensor Inputs

Sensor Variables	**Function**	**Description**
$V_1 \ldots V_4$	Object detection	Robot, block, no obstacle
G_1	Gripper status	Holding block, no block
S_1	Sign detection	Red, blue, orange, pink, green
D_1	Heading	North, east, west, south

Deciphering signs relative to the robot's current frame of reference makes these tasks particularly difficult given a fitness function that measures success in terms of reaching the goal location. The two versions of the task considered here are (1) where the controller has access to a compass sensor at each time-step and (2) where compass sensor readings are penalized or restricted. We shall refer to

the former variant as *compass-enabled* and the latter as *compass-restricted*. Even the simpler compass-enabled version is found to be unsolvable for predetermined fixed-network topologies that lack regulation (see Results and Discussion).

Table 2. Basis Behaviors

Order	Behavior	Description
1	Pick-Up/Put-Down	Pick up or put down obstacle
2	Move forward	Move one square forward
3	Turn right	Turn 90° right
4	Turn left	Turn 90° left
5, 7, 9, 11	Bit set[1]	Set memory bit i to 1, $i = 1 \cdots 4$
6, 8, 10, 12	Bit clear[1]	Set memory bit i to 0, $i = 1 \cdots 4$
13	Get hint[2]	Get current heading (D_1)

[1] Behaviors disabled for recurrent and memory neuron architectures
[2] Behaviors disabled under certain conditions (see text)

In the compass-restricted version, the controller must perform sign following knowing just its initial heading thus requiring the controller to predict and keep track of the robot heading (ego-orientation) in addition to accomplishing the other subtasks described earlier. Keeping track of long term dependencies is acknowledged to be difficult with recurrent connections [7] making the sign-following task a good benchmark for multistate architectures. The robot in this case has access to one additional behavior, the 'get hint' behavior, which interrogates the compass for the 'true' heading. However, the fitness function incrementally penalizes and restricts the number of hints used. The fitness function for a given run is defined as

$$f_i = \begin{cases} \dfrac{1}{1 + \beta n_{\text{hint}}/16}, & \text{if goal is reached} \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

The total fitness function is averaged over all runs. For the compass-enabled variant, $\beta = 0$ although the robot always knows the compass direction. So when the goal is achieved $f_i = 1$; otherwise, $f_i = 0$. For the compass-restricted one the reward for success is discounted according to the number n_{hint} of hints that have been used. However, for the first 5,000 generations, the robot is not penalized for using hints; hence $\beta = 0$. For the subsequent 10,000 generations, $\beta = 1$ but the hint can only be used in the first four time-steps. (This allows the robot to get the true direction reading as it starts out.) After 15,000 generations, hints are proscribed altogether. These parameters were found through experimentation to work well for this task.

The evolutionary performance of various control system architectures is compared for the two variants of the sign-following task (see Fig. 6). The robot's world is a 20×20 grid with 80 uniformly distributed obstacles and 40 randomly distributed mines (except along the path to the goal). The fitness is averaged over 100 runs with different initial conditions, the elapsed time for each run being limited to 100 timesteps. Controllers that lack recurrent connections or memory neurons have access to four memory bits, which can be manipulated

using the defined basis behaviors. Evolution assumes a population of 100 individuals in each generation and a tournament size of 24. The crossover probability $p_c = 0.7$, the mutation probability $p_m = 0.005$ and the transcription error rate $p_{te} = 0.005$.

5 Results and Discussion

Fig. 6a shows the average (population best) fitness of the controllers evaluated in each generation for the compass-enabled variant of the sign-following task. For some comparison, a fixed-topology recurrent network with 9 hidden and 4 output neurons is also shown. (Although this is typical of the results obtained for such a network, we did not optimize performance in any way.) Fixed-topology networks tend to have more 'active' synaptic connections present (all neurons are active) and thus more spurious neurons need to be dealt with simultaneously. The ANT topology with the coarse-coding neuroregulatory mechanism disabled (using modular activation function) shows better performance than fixed topologies (including H-ESP [7]) but not sufficient to complete the task for all the initial conditions tested.

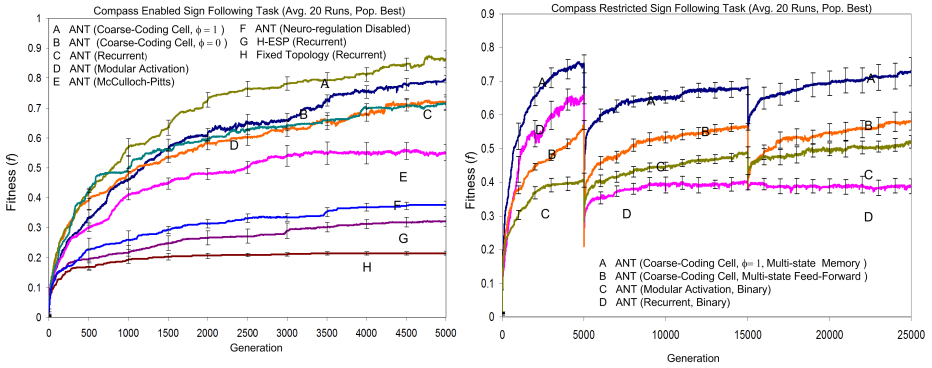

Fig. 6. Comparison of performance for (a) compass-enabled task and (b) compass-restricted task.

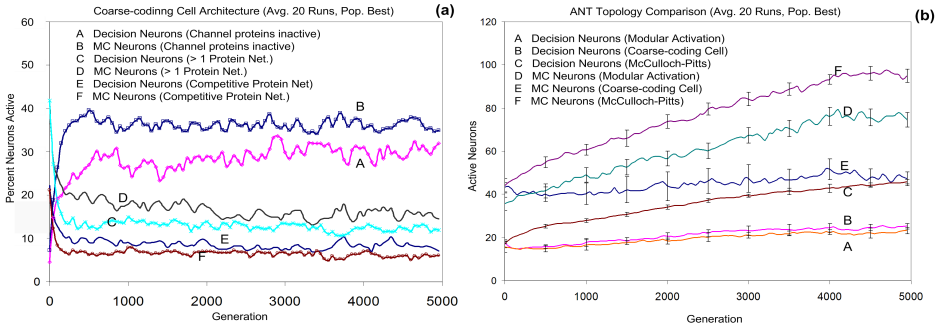

Fig. 7. Comparison of (a) coarse-coding cell models and (b) ANT topologies.

In contrast, heterogeneous architectures, whether predefined (modular activation function) or developmental, outperform other architectures. In the coarse-coding cell

model, selection is favoring increased complexity among decision neurons (Fig. 7a) but this is followed by a gradual simplification in neuron structure. This simplification is possible with all the channel proteins disabled, resulting in constant output (independent of sensory input).

While the modular activation function also allows for heterogeneity, it is a fixed-cell architecture that does not facilitate intracellular competition nor 'complexification.' Overall, the 'expressed' structure of the coarse-coding cell model is using fewer tuning parameters (both weights and thresholds) than with the modular activation function as it converges to a solution (Fig. 7b).

Fig. 7b also shows that many messenger-channels protein networks have redundant action proteins implying a 'weak' cooperative setup. A weak cooperative setup is advantageous, allowing for cooperative and competitive tendencies and can better facilitate a transition between the two. The added benefit of the heterogeneous coarse-coding cell model is that it also facilities memory functionality through gating. The evolutionary performance with gated-memory functionality ($\phi = 1$) shows a definite improvement over recurrent architectures for both tasks (Fig. 6a).

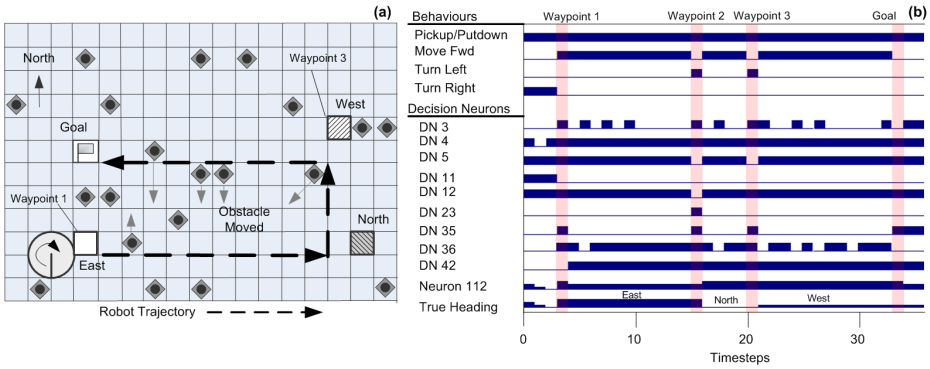

Fig. 8.(a) Typical ANT solution (fitness $f = 0.99$) using multistate neurons (memory model) for the compass-restricted sign-following task. (b) Output behavior and function of memory neurons.

For the compass-restricted version of the task (Fig. 6b), only the multistate coarse-coding cell models can reach a fitness of 0.9 within 25,000 generations at least once. The binary-state neuron model with feedforward connections (and access to 4 memory bits) showed comparable results with a fitness of 0.8 at least once while the binary recurrent architecture performs poorly. Among the four architectures, the binary recurrent model (localized memory) lacks built-in competitive network mechanisms for memory representation. For the multistate model, these competitive mechanisms are the protein networks resident in each cell, while similar functionality is present within the feedforward architecture, where neuron ensembles compete for control of the 4 memory bits (see [17]).

These results also show *intra*cellular competition has an advantage over *inter*cellular competition for memory representation. The multistate neuron model with memory functionality outperformed the multistate feedforward model (Fig 6b). It should also be noted that the model is more favorable in helping to map expected compass sensory input due to the inclusion of increment/decrement and reset 'actions' (see Fig. 8).

6 Conclusions

A developmental coarse-coding neuron model that facilitates heterogeneity and cell extensibility is found to exploit gated memory functionality. This architecture is found to be more effective than recurrent architectures for two memory -dependent variants of the unlabeled sign-following robotic task. The neuron model produces smaller networks on average than a predefined heterogeneous cell model and is driven by competitive/cooperative dynamic leading to specialization.

ACKNOWLEDGEMENTS. The authors would like to thank the reviewers for their helpful comments.

References

1. Albus, J. S., A theory of cerebellar function. *Mathematical Biosciences*, (1971)
2. Astor J.C., Adami C., A developmental model for the evolution of artificial neural networks. *ALife* **6**(3) (2000) 189-218.
3. Edelman G., *Neural Darwinism. The Theory of Neuronal Group Selection*, Basic Books, New York (1987).
4. Eggenberger, P., Gmez G., and Pfeifer R., Evolving the morphology of a neural net for controlling a foveating retina, *ALife*, **8** (2002) 243-251.
5. Eggenberger, P., Evolving Morphologies of simulated 3d organism based on Differential Gene Expression, Proc. 4th European Conference on ALife (1997).
6. Graves, A., Beringer, N., Schmidhuber, J., A Comparison Between Spiking and Differentiable Recurrent Neural Networks on Spoken Digit Recognition, Proc. of the 2nd IASTED Int. Conf. on Neural Nets, NCI (2004).
7. Gomez, F., Schmidhuber, J., CoEvolving Recurrent Neurons Learn Deep Memory POMDPs, GECCO (2005) 491-498.
8. Gruau, F., Automatic definition of modular neural networks. *Adaptive Behaviours*, **3** (1994) 151–183.
9. Hochreiter, S., and Schmidhuber, J., Long Short-Term Memory, *Neural Computation*, **9**(8)(1997) 1735–1780.
10. Husbands, P., Evolving Robot Behaviours With Diffusing Gas Networks, EvoRobots (1998) 71–86.
11. Sims K., Evolving 3D Morphology and Behavior by Competition, Proc. of Artificial Life IV, MIT Press, (1994) 28–39.
12. Kirschner, M., Gerhart, J., Evolvability, *Proc. Natl. Acad. Sci.* (PNAS), (1998)
13. Nolfi, S., Floreano D., *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press (2000) 13–15.
14. Roggen D., Floreano, D., Mattiussi, C., A Morphogenetic Evolutionary System: Phylogenesis of the POEtic Tissue, Int. Conf on Evolvable Systems (2003) 153–164.
15. Roggen D., Federici D., Multi-cellular Development: Is There Scalability and Robustnes to Gain?, Proc. of Parallel Problem Solving from Nature (2004) 391–400.
16. Thangavelautham, J., DEleuterio G. M. T., A Coarse-Coding Framework for a Gene-Regulatory-Based Artificial Neural Tissue, Advances In Artificial Life: Proc. of the 8th European Conf. on ALife (2005).
17. Thangavelautham, J., Alexander S., Boucher D., Richard J., D'Eleuterio G.M.T, Evolving a Scalable Multirobot Controller Using an Artificial Neural Tissue Paradigm, IEEE ICRA (2007).
18. Thangavelautham, J., D'Eleuterio, G.M.T, A Neuroevolutionary Approach to Emergent Task Decomposition, Proc. of 8th PPSN (2004) 991–1000.