

Supplementing Evolutionary Developmental Systems with Abstract Models of Neurogenesis

Keith L. Downing
Complex Adaptive Organically-Inspired Systems (CAOS) Group
Department of Computer Science
The Norwegian University of Science and Technology
7491 Trondheim, Norway
keithd@idi.ntnu.no

ABSTRACT

This work investigates an intermediate abstraction level, that of neural groups, for modelling the development of complex artificial neural networks. Based on Neural Darwinism [5], Displacement Theory [4] and The Neuromeric Model [18], our DEACANN system avoids the complexities of axonal and dendritic growth while maintaining key aspects of cell signalling, competition and cooperation that appear to govern the formation of neural topologies in nature. DEACANN also includes a genetic-algorithm for evolving developmental recipes, and the mature networks employ several forms of learning.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning—*Connectionism and Neural Nets*

General Terms

Algorithms

Keywords

Genetic Algorithms, EvoDevo, Neural Networks, Displacement Theory

1. INTRODUCTION

Nervous systems of higher vertebrates have a clearly hierarchical structure, with sensory inputs translated into motor outputs in a fast, purely reactive manner at the lowest levels, but with these same inputs propagating to higher neural structures, whose delayed effects upon motor activity reflect more advanced processes such as sensory integration, memory retrieval, prediction and even planning. However, bi-directional signalling is extremely prevalent among cerebral regions, indicating more of a heterarchical than purely hierarchical organization, as shown in Figure 1.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

Many modern Artificial Intelligence (AI) systems that attempt to mimic aspects of higher intelligence exploit similar heterarchical organization. Unfortunately, the closest contemporary computational analog to cerebral mechanisms, artificial neural networks (ANNs), are difficult to train on heterarchical structures. Many researchers have successfully utilized evolutionary algorithms (EAs) to evolve weight vectors for hierarchical and heterarchical ANNs [12, 21, 14], but these direct codings (one gene per weight) do not scale well to the sizes of networks needed for complex, beyond-reactive tasks.

Developmental approaches, in which EA genomes specify a recipe, not a blueprint, for ANN formation, arose to combat the scalability problem [10, 2, 8, 20]. In a seminal article on developmental EAs (DEAs) [17], Stanley and Miikkulainen attack the standard classification of DEAs as too superficial and implementation specific. As an alternative, they propose a biologically-based set of dimensions for distinguishing DEAs. Although somewhat orthogonal to the classic developmental processes cited by Wolpert [19] (cleavage division, pattern formation, morphogenesis, differentiation and growth), Stanley and Miikkulainen's dimensions (cell fate, targeting, heterochrony, canalization and complexification) also focus on the cellular (and even genetic) levels.

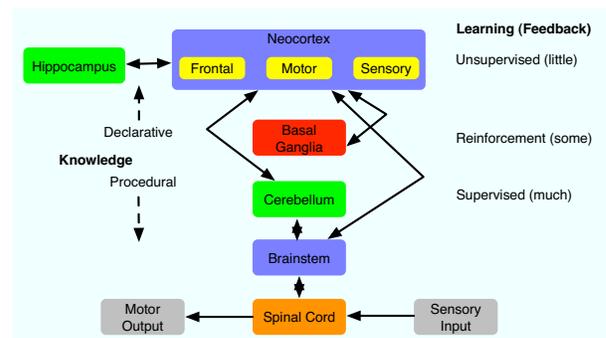


Figure 1: Mammalian sensorimotor control hierarchy

Unfortunately, the distance between these low levels and the control heterarchies seen in the brain is great. Even incremental complexification approaches, as in NEAT [9], cannot achieve the multi-level neural topologies seen in nature. Although Stanley and Miikkulainen argue for abstractions

well above the level of cellular migration, chemical signalling and neuritic growth, this paper argues for yet another step up from the cellular level.

In [17], Stanley and Miikkulainen discuss heterochrony: the effects of timing upon developing phenotypes. The strong adaptability of developing embryos normally prevents timing changes from being fatal, but instead, allows greater exploration of phenotypic space via small genotypic (timing) variations. Exploration via heterochrony is reasonably safe during early and late development, when intercellular communication is relatively low. However, the middle, *phylotypic* stage is much less flexible due to a high level of global chemical signalling.

Interestingly enough, this phylotypic stage is precisely the stage during which the embryos of many different species look alike. In fact, by examining neural structures during the phylotypic stage, one finds structural similarities that can form the basis for a general model of the development of heterarchical ANNs.

This article examines one of the most popular (among developmental neuroscientists) characterizations of the phylotypic stage, the Neuromeric Model, and describes a DEA based upon it.

2. GROUP-LEVEL PRINCIPLES OF NEUROGENESIS

In *Principles of Brain Evolution* [18], Striedter, a comparative neurobiologist, reviews a host of useful principles, at many abstraction levels, for understanding neurogenesis, neuroevolution, and their interaction. Two of the key principles are the Neuromeric Model and Displacement Theory. The former addresses the spatial arrangement of brain regions, while the latter explains the topology of connections between them. Together, they provide a promising intermediate level of abstraction for DEAs that grow artificial neural networks.

2.1 The Neuromeric Model

In 1953, Bergquist and Kallen [3] noticed that all vertebrate embryos have a similar elongated, segmented hind-brain during the phylotypic stage. The ringed segments, termed *neuromeres*, are modular zones of high cell division (to form neurons) and radial migration (of neurons to their proper layer). Later, Puelles and Rubenstein [16] found that this pattern encompassed the midbrain and forebrain as well. They also provided genetic evidence that Hox and Hox-like genes control this segmentation, just as they control the subdivisions of the body's central axis. Hence, this revised Neuromeric Model views the entire developing brain as an elongated series of ringed modules, within which develop layers of neuron cell bodies (i.e., gray matter).

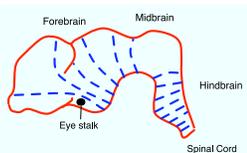


Figure 2: The contemporary neuromeric model.

Figure 2 sketches the basic neuromeric structure of the vertebrate phylotype. The hindbrain neuromeres develop

into brain regions such as the cerebellum and pons, which are tightly tied to sensory and motor systems, while the mid-brain and forebrain segments become areas such as the basal ganglia, hippocampus and prefrontal cortex, all of which are involved in high-level cognitive processes. Hence, the Neuromeric Model provides the perfect developmental scaffolding for the control heterarchy of Figure 1.

2.2 Displacement Theory

The Neuromeric Model provides a useful structural bias for generating neuron sub-populations in a DEA for ANNs, while abstracting away many cellular-level details. Deacon's [4] Displacement Theory (DT) compliments the Neuromeric Model by explaining the interactions between the sizes of neural sub-populations and their inter-connectivity, again, while remaining above the cellular level of abstraction.

The basis of DT lies in Edelman's [5] Darwinistic view of neurogenesis, known as The Theory of Neural Group Selection (TNGS). In this view, neurons undergo a selective process wherein only those that grow axons to, and receive axons from, other neurons will reach maturity. Essentially, neurons are involved in a *survival of the best networkers* competition. DT expounds on TNGS by proposing that the networking competition during early development enables brains to *scale to fit* the body's sensory and motor apparatus. In short, primary sensory and motor areas of the brain are sized according to their immediate inputs or outputs, respectively. Secondary region sizes derive from those of the primary regions, and deep cortical structures grow or shrink to fit their input and output sources.

Figure 3 conveys the essence of TNGS and DT. Note a) the expansion of the 3 neuron groups along the path from the largest sensory input, S1, to the largest motor output, M2, and b) the decline of groups B and D, which lose the competition for C's dendrites and C's axons, respectively.

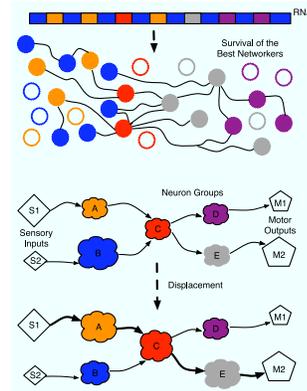


Figure 3: (Above) The Theory of Neural Group Selection (TNGS), wherein neurons survive only if they establish efferent and afferent connections. Dying neurons are unfilled in the figure. (Below) Displacement Theory (DT), in which neuron groups with good networking possibilities expand, while others shrink. Sizes of sensory, motor, and neuron group icons depict sizes of the corresponding pools.

Developmental neuroscience clearly supports and employs the key tenets of DT. For example, Fuster [7] documents the

earlier maturation of posterior brain regions (used in early sensory processing) and late maturation of frontal areas such as the prefrontal cortex (PFC). Striedter [18] combines DT with the work of Finlay and Darlington [6] to explain the trend of greater neocortical (and especially PFC) control of lower brain regions in higher organisms. First, Finlay and Darlington show that *late equals large* in neurogenesis: larger brain regions are those that mature later in development. Second, a key corollary of Deacon’s DT is that *large equals well-connected*: big brain regions send many axons to other regions and thereby have a significant amount of control over them. Together, these show how small changes in developmental timing (in higher mammals) have enabled the frontal regions to mature later, hence grow larger, and hence exhibit greater control over a wide variety of cortical areas. And greater frontal control correlates well with behavioral sophistication, as illustrated by Nakajima et al.’s [13] comparisons of manual dexterity vis-a-vis frontal control of motor areas in mammals such as cats, monkeys and humans.

Together, these theories paint neurogenesis as a self-organizing process in which genetics determines the neuromeric structure, the basic properties of neurons in different layers of the neuromeres, and the maturation speed of neural regions, but the final sizes of these regions and their interconnectivity stem from self-organizing processes wherein neuronal subpopulations essentially compete for the right to cooperate (via synaptic signalling) with other subpopulations.

3. GROUP-LEVEL DEVELOPMENT IN DEACANN

Designed to emulate the essential elements of neuromere formation, Neural Group Selection and Displacement Theory, while avoiding the computationally intensive simulation of axon growth, the developmental algorithm employed in DEACANN involves three phases.

In phase I, **Translation**, the binary genome is converted into a set of neuromeres, each containing one or more neuron groups, whose basic properties are genetically determined.

In phase II, **Displacement**, the sizes of each neuron group and intra- and inter-group invasion strengths undergo repeated cycles of modification. Since (in both TNGS and DT) group size affects invasion strength, and vice versa, several iterations are required to translate the initial group masks, inter-group distances, growth limits and proximities (to sensory inputs and motor outputs) into group sizes and relative degrees of inter-group connectivity.

In phase III, **Instantiation**, the final group sizes and invasion strengths are used to generate a) populations of neurons for each group, and b) connections between individual neurons in the same and different groups. Although not detailed in this paper, some of these connections specify neuromodulatory signalling pathways, wherein single neurons send messages to entire neuron groups.

3.1 Translation

The DEACANN genome needs to support the development of neuromere chains, where each segment may consist of several distinct neuron groups. All neurons in a group should be similar, and a group type may repeat several times in the same or in different neuromeres. For example, a layer

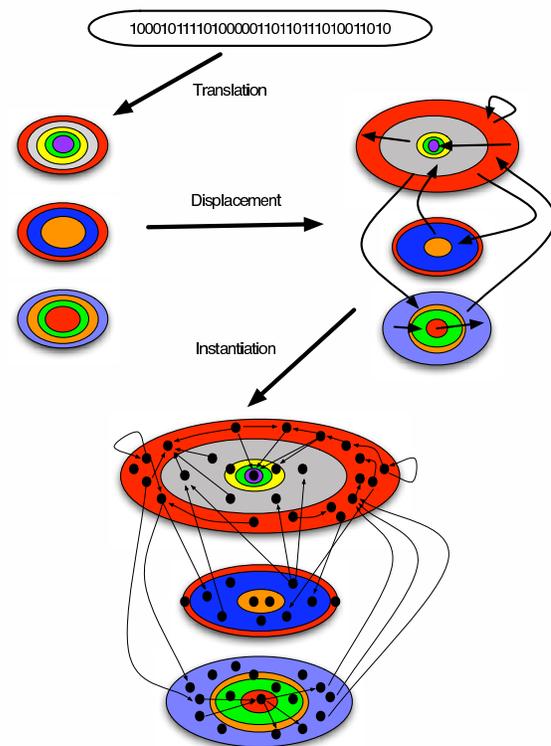


Figure 4: DEACANN’s 3-stage developmental model. Each circle denotes a neuromere, while each concentric ring represents a neuron group, with ring width proportional to group size. In the upper right, arrows between neuron groups indicate biases (quantified as invasion factors - see below) that generally lead to axonal connections during the instantiation phase. In the bottom section, small dots are neurons, and their axons are thin, outward-bound arrows.

of lateral inhibitors may occur several places in the developed neural network.

Although one could rely on evolution to duplicate the genes that encode these groups, the probability of such copying (in a linear genome) decreases dramatically as gene complexity rises. Although genetic programming with sub-routines [11] provides one method of modularizing reusable genes, the computational overhead of GP seems unnecessary in this case, given the objective of evolving sets of neuron-group parameters, not actual growth procedures.

In this special case, where neither the number of neuromeres nor the number of groups per neuromere is fixed, we can achieve modularity and reuse within a linear GA chromosome by simply removing the standard GA constraint of fixed gene locations and adopting a tag-based addressing system, as shown in Figure 5.

In this indirect representation, binary tags denote the start of a neuromere specification. In the example of Figure 5, this tag is 11111, with length $k = 5$. The user can specify the degree to which any k -length segment of the chromosome must match this tag to qualify as a hit. A matching segment is called a *neuromere header*. Here, we assume a match de-

gree of 100 % for neuromere headers, i.e., they must match the neuromere tag exactly.

Once a neuromere header is found, the bits directly following it are interpreted as the *neuromere specification*. In the current version of DEACANN, this specification is simply a single value denoting a neuron-group tag. In Figure 5, the single neuromere specification (denoted by the pentagon labelled "N") has 01010 following its header. Thus, the neuron-group tag for this neuromere is 01010.

The translator then scans the (entire) chromosome in search of segments that match the neuron-group tag, 01010. In this example, we assume that a 80% or more of a chromosome segment's bits must match the tag. Each such matching segment denotes the header of a neuron-group specification, and the bits following the header are translated into the various parameters of the group, such as axon and dendrite masks, learning rates, etc. (which are detailed below).

Thus, for each neuromere specification header that is found, a neuron group tag is read from the chromosome and used as the basis for a complete scan of the chromosome for the specifications of the neuron groups of that neuromere. During a scan, the group specifications cannot overlap. However, the group specification for one neuromere may be ignored or interpreted differently during the scan for another neuromere's groups. Similarly, the (rather short) neuromere specifications cannot overlap.

A match degree less than 100% allows neuromeres to share some, but not necessarily all, group specifications, since each neuromere has a potentially unique group tag. For example, assuming a common group-tag-match degree of 75%, if neuromere N1 has group tag 1111 and neuromere N2 has group tag 1010, then both neuromeres would share any group specification with header 1011 or 1110, but only N1 would include groups with header 0111 and only N2 would select groups headed with 1000.

This representation supports complexification via the well-known combination of genetic duplication and differentiation [1]. Being relatively short, neuromere coding regions are frequently copied in their entirety, thus creating additional neuromeres in later generations. Any mutation to the new neuromere's group tag then provides the potential for differentiation, since this enables the new neuromere to inherit some, but not all, of the neuron-group types from its ancestor neuromere.

The neuron-group specification consists of a contiguous string of bits that encode the following parameters:

1. Axon mask – Masks are abstractions of cadherins, ephrins, [19] and other chemicals that govern the attraction and repulsion of axons during their migration toward dendritic targets. In general, the attractiveness of one neuron group for another is directly proportional to the complementarity of the axon and dendrite masks in the two groups, and inversely proportional to the distance between the groups.
2. Dendrite mask
3. Axon sharing – The sharing parameters indirectly control the patterns of connectivity formed between two groups during the Instantiation phase of development.
4. Dendrite sharing
5. Postsynaptic effect – Indicates whether group neurons have excitatory or inhibitory effects upon targets.

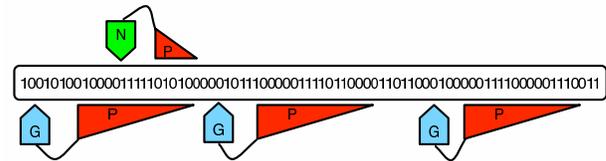


Figure 5: Translation of a bit string into the parameters for the neuron groups of a neuromere. Pentagons point to the start of neuromere (N) and group (G) tags, while triangles denote the extent of the parameter bits (P) that follow each tag. In this example, the tag for neuromeres is 11111 and must be matched exactly, while group tags must match on at least 4 of their 5 bits. The single parameter for the neuromere is the tag for its groups, 01010, which appears under the triangle attached to pentagon N. Hence, all strings that match 80% of 01010 mark the start of group specifications.

6. Neuromodulator sent – Governs whether group neurons send neuromodulators (and of what type) instead of normal action potentials when excited. If the group is excitatory (inhibitory), then it's neuromodulator will always have an excitatory (inhibitory) effect upon targets.
7. Neuromodulator received – Neuromodulator (if any) to which the group neurons are sensitive.
8. Growth limit – Number of developmental rounds in which the group neurons will participate.
9. Learning rule – One of several Hebbian and anti-Hebbian synaptic-modification schemes, with the same rule applied to all incoming connections of a group's neurons.
10. Learning rate – Degree of synaptic-strength change during learning-rule application.

3.2 Displacement

Phase II of development is the heart of the DEACANN approach. It simulates the interaction between neuron groups, both within and between neuromeres, but without simulating the actual growth of axons and dendritic trees. Instead, three interacting factors are iteratively updated: 1) the cardinality of the neuron sets in each group, S_i , 2) the invasion strength of each group relative to itself and others, $I_{i,j}$, and 3) the connectivity of each neuron group, C_i .

Invasion strength, $I_{i,j}$, represents the propensity of neurons in group i for sending axons to invade targets in group j . It is a function of the distance between the two groups, $D_{i,j}$, the compatibility of the axon mask for i with the dendritic mask for j , $M_{i,j}$, and the sizes of groups i and j , S_i and S_j . The basic update formula is given in equation 1, in which α_m and α_d are weighting constants with typical values of 0.5 and 1, respectively, and T is the total elapsed time since the beginning of Phase II. The inclusion of T allows groups that are distant from one another but otherwise compatible, i.e. high $M_{i,j}$, to eventually hook up if both groups can maintain a reasonable size during the early stages of displacement.

$$I_{i,j} = S_i S_j (T + 1) \frac{\alpha_m M_{i,j}}{1 + \alpha_d D_{i,j}} \quad (1)$$

For each group, i , the total invasion force into, I_i^{in} , and out from i , I_i^{out} are derived from the pairwise invasion strengths as:

$$I_i^{in} = \sum_{j \in G} I_{j,i} \quad (2)$$

and

$$I_i^{out} = \sum_{j \in G} I_{i,j} \quad (3)$$

Neuron groups compete for targets via the following modification to invasion strength:

$$I_{i,j} \leftarrow (1 - \gamma) \frac{I_j^{in}}{N_g} + \gamma I_{i,j} \quad (4)$$

Here, $\frac{I_j^{in}}{N_g}$ is the average intensity of invasions into group j ; N_g is the total number of neuron groups in the ANN. γ denotes the competition intensity, a constant with a range of [0 1] and typical value of 0.5. Higher values of γ will penalize intensities that are below the average, thus squeezing out weak invaders. Conversely, a γ closer to 0 will push all invasion strengths closer to the average, thus more evenly distributing access to group j 's dendrites.

The competition-modified invasion strengths are re-summed to update the invasion forces, and the estimated - remember, no actual connections have yet been formed - connectivity of each neuron group is then computed as:

$$C_i = I_i^{in} + I_i^{out} \quad (5)$$

To reflect the basic tenet of Neural Darwinism - the best networkers proliferate, while poorly-connected neurons die - neuron-group sizes are updated as a function of connectivity:

$$\Delta S_i = \alpha_g (C_i - \bar{C}) \quad (6)$$

where \bar{C} is the average connectivity over all neuron groups, and α_g is a growth constant, typically 0.1. Notice that S_i does not occur in this update formula; its effect is already present in the contribution of I_i^{out} to C_i , since each addend of I_i^{out} involves S_i .

These updates of $I_{i,j}$, C_i and S_i are repeated a maximum of R (a user-defined parameter typically in the range [5 10]) developmental rounds. Any group with a growth limit, G_i less than R will not participate in the final $R-G_i$ rounds. This embodies the *late equals large* principle, since neuron groups with longer growth limits will tend to have higher connectivity - remember the affect of T in equation 1 - and hence larger sizes.

Initially, each group size is 1, and typically, each input and output neuron constitutes its own group. Inputs are assumed to enter through the outermost groups of neuromere 0, while outputs exit through its innermost groups. Groups that are spatially adjacent or otherwise attractive as targets for several input and/or output groups will have high incoming invasion forces, and thus high connectivity and increasing sizes, exactly as prescribed by Displacement Theory.

3.3 Instantiation

Given the S_i and $I_{i,j}$ values for each group and group pair, respectively, Phase III of development uses these values to bias the generation of neurons for each group and connections between them.

To generate neurons, all group sizes, except the input and output groups, are normalized to produce distribution fractions for a fixed total number of neurons, N .

Connections are formed by considering each group, i , and all invasion factors into that group. These factors are normalized to produce the values $I_{j,i}$ for all groups j . For each potentially invading group j , corresponding neurons in groups i and j are considered. For each pair, a connection from the j neuron to the i neuron is formed with probability $\tilde{I}_{j,i}$.

Once formed, a connection from the k th neuron of group j to the k th neuron of group i neuron can spawn further connections of a convergent, divergent or parallel form, depending upon the axonal sharing factor of group j , λ_j^{out} and the dendritic sharing factor of group i , λ_i^{in} . In the process described below, the k th neuron of group is denoted $G_i(k)$

- $k_i = k$, $k_j = k$
- Repeat
 - Create connection $G_j(k_j) \rightarrow G_i(k_i)$
 - Generate random fractions r_i and r_j
 - If ($r_i \leq \lambda_i^{in} \wedge r_j \leq \lambda_j^{out}$) then $k_i \leftarrow k_i + 1$ and $k_j \leftarrow k_j + 1$
 - Else if $r_i \leq \lambda_i^{in}$ then $k_i \leftarrow k_i + 1$
 - Else if $r_j \leq \lambda_j^{out}$ then $k_j \leftarrow k_j + 1$
- Until ($r_i > \lambda_i^{in} \wedge r_j > \lambda_j^{out}$)

Briefly, if both sharing factors are high, then the first conditional will often be true and both indices will be incremented, causing the next pair of corresponding neurons to be paired. Several rounds of this will create a set of parallel connections, similar to a topological map. If dendritic sharing is high but axonal sharing low, then the second, but not the first, conditional will often be true, thus causing the presynaptic neuron to remain fixed, while the postsynaptic neuron becomes the neighbor to $G_i(k_i)$. A few repeats of this situation forms a divergent connection pattern from group j to i . Conversely, if dendritic sharing is low but axonal sharing high, the third condition is frequently triggered and a convergent pattern results. Finally, if both sharing values are low, then the loop exits early and only one or a few connections are created. These outcomes are summarized in Figure 6.

This approach has a drawback with respect to evolvability: since synapses are generated stochastically, the same genome can give rise to different phenotype topologies. Hence, a good phenotype may disappear from the population if its genotype (G), and G 's near neighbors in genotype space, all develop into low-fitness phenotypes in any particular generation.

To reduce this non-determinism during development, an alternate approach, *scaled pairing*, begins by comparing $\tilde{I}_{j,i}$ to a threshold (0.2 in the results below). If it is above, then corresponding neurons in groups i and j are immediately coupled, with a *scaling factor* of $\tilde{I}_{j,i}$ attached to each

link. As detailed in [15], this fixed factor modifies the link’s weight (which, due to learning, is not fixed). The stochastic computation of divergent and convergent links can then proceed as above, with each new connection also receiving $\tilde{I}_{j,i}$ as a scaling factor.

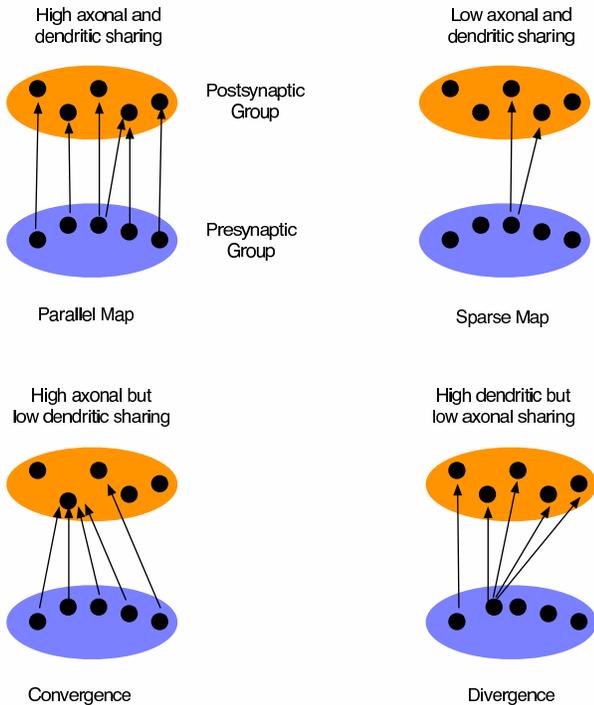


Figure 6: Inter-group connection topologies as a function of axonal and dendritic sharing factors.

4. RESULTS

Three-way adaptive (a.k.a. TRIDAP) systems that combine evolution, development and learning create a rather ominous search space, as work with DEACANN has clearly revealed. Several problems were unsuccessfully attempted, where *success* is defined very loosely as the ability to show improved fitness over the course of evolution. These included short-term-memory and simple tracking tasks.

However, a simple movement control problem, the *starfish*, does allow DEACANN to consistently show evolutionary progress. In this task, a radially-symmetric k-limbed animat must move as far away from its’ starting point as possible in a limited time. Each limb has two *muscles*, one for pulling in the limb’s direction, and one for pushing in the opposite direction. A motor neuron is devoted to each of these muscles, so a k-limbed starfish has 2k motor neurons. The firing rate of each neuron affects the strength with which its muscle contracts, and the magnitude and direction of starfish movement is then the resultant vector averaged over all the pulling and pushing contractions, weighted by their intensities. The four input neurons register positive and negative velocity components along the x and y axes.

To make the task more interesting, each muscle has a fatigue factor that increases (decreases) each time its’ motor neuron’s membrane potential is above (below) a firing

threshold, which is 0.25 in these examples. The fatigue factor inversely affects the ability of motor neuron’ activity to translate into muscle activity, with fully fatigued muscles unable to contract. In this model, muscles still contract, but with weak intensity, when the motor neuron is above zero but below the firing threshold.

When tested on a 5-limbed starfish for 50 timesteps, DEACANN typically evolves progressively better controllers; involving multiple interacting neuron groups and neuromeres that produce cooperative motor-neuron firing patterns. Although a starfish can move a maximum of one distance unit per timestep, fatigue prohibits a string of 50 such maximum-distance movements in the same direction, thus rendering the calculation of an optimal 50-step distance somewhat difficult. Running a population of 20 individuals for 20 generations in DEACANN typically yields a top performance of 14 to 16 distance units. These runs employ fitness-proportionate selection, with three-individual elitism and culling of the lower 25 % of each generation. Genomes are 200 bits long; development has a maximum of 5 rounds, employs scaled pairing and randomly initializes all weights to values in (0, 0.1]. An individual’s fitness is simply its’ distance from the start point after 50 time steps.

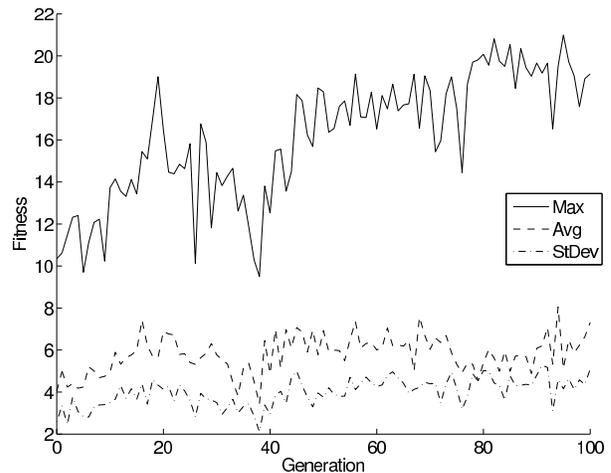


Figure 7: Fitness progression of a 20-individual DEACANN population on the 5-limbed starfish problem.

Figure 7 shows a standard evolutionary progression, while Figure 8 shows the best individual of generation 100, which moves nearly 20 units, while Figure 9 illustrates the winning strategy. A net pushing force is generated from arms 1 and 5, while a net pulling force occurs in arms 3 and 4. Arm 2 has competing pulling and pushing forces, which largely cancel one another. In limbs 1, 3, 4 and 5, a dominant force arises due to a fatigued opposing force. For example, in limb 1, the motor neuron for pushing fires only weakly, while that for pulling fires strongly. This causes the former to work weakly, but without accumulating fatigue, while the latter quickly fatigues; and since its’ motor neuron never rests, the limb-1 pulling muscle remains inactive for the remainder of the run. Since all motor neurons have the same default time constants, biases, etc., their different firing intensities stem from heterogeneous inputs and learning rules. As shown in Figure 8, 3 of the large intermediate neuromere’s 4 neuron

groups are inhibitory, so differential connectivity between the 10 motor neurons and the different layers of neuromere 1 (along with different learning rules) leads to heterogeneous combinations of inhibition and excitation, and thus to the special mixture of pushing and pulling intensities that facilitates 20-units of movement.

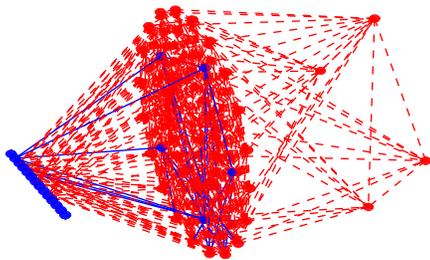


Figure 8: Neural network of the best-of-generation individual after 100 generations of a 20-individual DEACANN run on the 5-limbed starfish problem. Three neuromeres occupy successive slanted planes in the picture, and neuron groups are concentric rings within each plane. Darker neurons (with solid lines emanating from them) are excitatory, while the lighter circles (with dotted emerging lines) are inhibitory. The leftmost bar, neuromere 0, consists of the 4 input and 10 output neurons, each of which is its own neuron group.

As shown in Figure 7, even with elitism, DEACANN cannot always maintain its' maximum fitness across the generations. This stems from developmental stochasticity, with the main source of nondeterminism being the synaptic formation process (despite the use of scaled pairing), as discussed above.

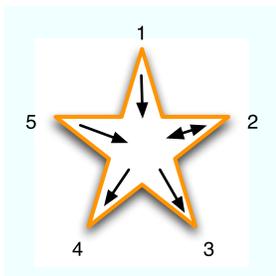


Figure 9: Main forces generated by the winning starfish of the 100-generation run.

5. CONCLUSION

Although DEACANN's intermediate level of abstraction saves considerable computational effort during development, it does involve a complex iterative process that is quadratic in the number of neuron groups. The neuron-group level a)

still affords the cooperative and competitive interactions at the heart of Neural Darwinism, and b) appears to be the proper level at which to explain the formation of large-scale brain topologies, via principles such as *late equals large* and *large equals well connected* and structures such as laminae and neuromeres [18]. In artificial TRIDAP systems, development can afford to work primarily at this intermediate level, since learning can (and should) handle the fine-tuning of individual synapses. Future work will investigate tasks in which learning plays a more obvious role than in the current starfish controllers.

6. REFERENCES

- [1] J. ALLMAN, *Evolving Brains*, W.H. Freeman and Company, New York, NY, 1999.
- [2] J. ASTOR AND C. ADAMI, *A developmental model for the evolution of artificial neural networks*, *Artificial Life*, 6 (2000), pp. 189–218.
- [3] H. BERGQUIST AND B. KALLEN, *On the development of neuromeres to migration areas in the vertebrate cerebral tube.*, *Acta Anat* (Basel), 18 (1953), pp. 65–63.
- [4] T. DEACON, *The Symbolic Species: The Co-evolution of Language and the Brain*, W.W. Norton and Company, New York, 1998.
- [5] G. EDELMAN AND G. TONONI, *A Universe of Consciousness*, Basic Books, New York, NY, 2000.
- [6] B. FINLAY AND R. DARLINGTON, *Linked regularities in the development and evolution of mammalian brains*, *Science*, 268 (1995), pp. 1578–1584.
- [7] J. FUSTER, *Cortex and Mind: Unifying Cognition*, Oxford University Press, Oxford, 2003.
- [8] F. GRUAU, *Genetic micro programming of neural networks*, in *Advances in Genetic Programming*, J. Kenneth E. Kinneer, ed., MIT Press, 1994, ch. 24, pp. 495–518.
- [9] KENNETH AND R. MIKKULAINEN, *Evolving neural networks through augmenting topologies*, *Evolutionary Computation*, 10 (2002), pp. 99–127.
- [10] H. KITANO, *Designing neural networks using genetic algorithms with graph generation system*, *Complex Systems*, 4 (1990), pp. 461–476.
- [11] J. R. KOZA, *Genetic Programming II: Automatic Discovery of Reusable Programs*, MIT Press, Cambridge, MA, 1994.
- [12] D. J. MONTANA AND L. D. DAVIS, *Training feedforward networks using genetic algorithms*, *Proceedings the Eleventh International Joint Conference on Artificial Intelligence*, (1989), pp. 762–767.
- [13] K. NAKAJIMA, M. A. MAIER, P. KIRKWOOD, AND R. LEMON, *Striking differences in transmission of corticospinal excitation to upper limb motoneurons in two primate species*, *Journal of Neurophysiology*, 84 (2000), pp. 698–709.
- [14] S. NOLFI AND D. FLOREANO, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*, The MIT Press, Cambridge, MA, 2000.
- [15] R. C. O'REILLY AND Y. MUNAKATA, *Computational Explorations in Cognitive Neuroscience*, The MIT Press, Cambridge, Massachusetts, 2000.
- [16] L. PUELLES AND J. RUBENSTEIN, *Expression patterns of homeobox and other putative regulatory genes in the embryonic mouse forebrain suggest a neuromeric organization.*, *Trends in Neuroscience*, 16 (1993), pp. 472–479.
- [17] K. STANLEY AND R. MIKKULAINEN, *A taxonomy for artificial embryogeny*, *Artificial Life*, 9 (2003), pp. 93–130.
- [18] G. F. STRIEDTER, *Principles of Brain Evolution*, Sinauer Associates, Sunderland, Massachusetts, 2005.
- [19] L. WOLPERT, *Principles of Development*, Oxford University Press, New York, 2002.
- [20] L. YAEGER, *Computational genetics, physiology, metabolism, neural systems, learning, vision and behavior or polyworld: Life in a new context*, in *Artificial Life III, Proceedings Volume XVII*, C. G. Langton, ed., Reading, Massachusetts, 1994, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, pp. 263–298.
- [21] X. YAO, *Evolving artificial neural networks*, *Proceedings of the IEEE*, 87 (1999), pp. 1423–1447.