Evolving Functional Symmetry in a Three Dimensional Model of an Elongated Organism

Ben Jones[1], Yaochu Jin[2], Bernhard Sendhoff[2], Xin Yao[1]
[1]School of Computer Science, The University of Birmingham, UK
[2]Honda Research Institute Europe, Offenbach, DE
B.H.Jones@cs.bham.ac.uk

Abstract

In evolutionary–developmental biology, it is well established that neural organization is coupled to a given organism's body-plan. Many theories attempt to underpin this coupling and the transitions involved during the organism's evolution, for example the transition from radial to bilateral symmetry. Before theoretically tackling these transitions however, we felt it essential to first address, in this paper, precisely why bilateral symmetry might be advantageous for a simple eel-like agent. We find that neural architectures affording the best motor-coordinated behavior (architectures that allow directional swimming of the agent), will readily emerge in a way that is functionally–bilaterally symmetric, suggesting therefore, that bilaterally symmetrical emergence for a long elongated creature can be essential if it needs to travel over some distance.

Introduction

The symmetrical properties of animals are mixed and varied. Typically, most higher organisms are bilaterally symmetric, that is to say, they can be partitioned into both dorsal and ventral halves. By comparison, more primitive organisms are radially symmetric and it is conjectured that the bilateral properties of higher organisms evolved from such radiata – and *both* from a common ancestor – during a process of *symmetry breaking* (e.g., Meinhardt (2002)). The general consensus is that the nervous systems of said organisms evolved in a coupled fashion so that they followed suit from body-plan architectural changes. As two fundamentally different examples, both the jellyfish (a radial organism) and the flatworm (a bilateral organism) demonstrate this principle in that their nervous system architectures have clearly evolved to reflect their body-plan morphologies.

Symmetry breaking is the evolutionary process that underlies the aforementioned change in body-plan symmetry. As discussed, this change is thought to have begun with a radial ancestor. Meinhardt (2002) considers gene homology as indicative of this common ancestry. More radical is the view that bilateral organization came about when a colony of individual polyps with Cnidarian (jellyfish) characteristics came together, see e.g., Collins and Valentine (2001); Holland (2003)). Further is the *Polyp with a half nerve net* scenario, attributed to Lacalli (1996). This argues that at some point during evolutionary history, a polyp started to crawl on its side, resulting in a build-up of the nervous system tissue in its ventral half and a concordant depletion in its dorsal half.

We will pick up on the issue of symmetry and although we do not account for the above theories, we will describe a very simple framework for testing the advantage of bilateral symmetry and the associated neural network (as a model nervous system) that emerges with this advantage (if indeed there is any advantage). On the one hand, we see this as a step in determining precisely why evolution favors particular bilateral body-plan nervous system couplings. If the above theories which all inherently argue that bilateral symmetry is evolutionarily advantageous, then we should hopefully observe its advantage for a simple agent in a simple environment. On the other hand, we are interested in how information processing might be structured in novel ways. We see our approach as one that enables us to study the coupling of neural architecture to body-plan symmetry. Although we do not strictly evolve the body-plan, we can still change the symmetry for a hypothetical system of muscles; and, since we fix the locations of these muscles around particular parts of the body, they can be considered as being part of the body-plan. Accordingly, if all muscles around the model organism are evolved to play a part in movement, then we will be able to partition the body-plan into several planes of symmetry and the muscle configuration can be regarded as being radially symmetric; whilst if only opposite muscles (those on the dorsal or ventral, or left or right parts of the animat) are evolved for movement, then we can hypothetically 'cut' the agent into halves and the muscle configuration can be regarded as being bilaterally symmetric. These symmetrical properties are not pre-defined, but will rather emerge if there is any evolutionary advantage.

We are not the first to investigate the coupling between body-plan morphology and neural network controllers. There are generally two bodies of researchers that have made related investigations. The first body is inter-

ested in modeling two dimensional models of undulatory organisms with a view to establishing some kind of undulatory behavior, and the type of neural controller that can bring this about; see for instance Zheng et al. (2004), Ekeberg (1993a,b), Sfakiotakis and Tsakiris (2006), Beauregard and Kennedy (2006), Ijspeert and Kodjabachian (1999). All of these studies share the common aim of understanding locomotion from a neuroscientific perspective. The second body of researchers are less interested in neuroscience but more interested in the behavior that can be evolved. Within this body, Karl Sims is one of the earliest proponents (Sims (1994a,b)) and many others have followed suit (Eggenberger (1997); Bongard and Paul (2000); also see Taylor and Massey (2001) for an extensive review). Most of these models are three dimensional and are implemented in part with powerful graphics libraries so as to provide the required visualizations and physical embodiments.

In terms of investigating body-plan symmetry, Bongard and Paul (2000) find that more locomotively efficient agents have a tendency towards evolving bilateral symmetry. In their model, they embed the neural controller into the agent's morphology so that both co-evolve. They further forgo any developmental process since they argue that one could inherently introduce symmetry; instead, they explicitly map the neuron weights and connectivities directly. However by doing so, the synaptic strengths and interconnectivities are de-coupled which in turn constrains the overall importance of neural network morphology. We argue that encoding a network at a greater level of morphological detail, so that neuron positional information has an actual bearing on connection strength, is essential, if we are to later observe any tendencies for different neurons to aggregate together and therefore potentially demonstrate central nervous system type characteristics. The 'GasNets' developed by Husbands et al. (1998) utilise similar neuron spatial information during a process of 'gas' diffusion, in which diffusing gasses play a crucial role in neuromodulation. In comparison, our model uses spatial information to determine connection strength rather than neuromodulatory signal effect.

Of further note is the work of Downing (2007) who constructed an evolutionary–developmental model of neurogenesis to bring about directional movement for a radially symmetric five-limbed 'starfish'. The model did not indicate how neural architecture may actually be coupled to body-plan morphology however.

Our own model has been constructed to meet our aim of investigating nervous system architecture/body-plan morphology couplings. In our simulations, both of these aspects co-evolve. Our motivation for this undertaking, is to initially elucidate the 'how' of this process (beginning with this paper), and our long-term goal is to better understand the 'why'. Thus we are interested in both information processing and the underlying evolutionary process. The model's task – that of directional swimming for an eel-like agent – is one of the simplest we could think of, yet it is also highly specialised inevitably requiring very task-specific couplings. This makes the problem non-trivial.

The rest of this paper is laid out as follows. We first outline and discuss some previous models of undulatory organisms. We secondly explain our model in more detail. Thirdly, we discuss our experimental results. We finally conclude this paper.

A Model of Undulatory Locomotion

Undulatory locomotion is a type of locomotion often employed by bilaterally symmetric creatures requiring directional movement (e.g. an eel). Refer to Gillis (1996) for a description of the underlying physics. Models of this type of behavior often adopt a spring mass damper system so that the mechanics are fluid and life-like. They secondly incorporate a friction model so that the modelled organism can actually move within its simulated world. Thirdly, they usually have a control mechanism, for example, a continuous time recurrent neural network (CTRNN). A CTRNN is often employed, because it is capable of exhibiting the *central pattern generating* dynamics that are essential for coordinated movement. A central pattern generator is a type of neural network that can by the very nature of its inherent dynamics, generate patterns of activity without any external input. For an extensive exploration of CTRNN dynamics, see Psujek et al. (2006); Beer and Gallagher (1992).

One of the earliest models attempting to use central pattern generators (CPGs) to model undulatory locomotion is that by Ekeberg (1993a,b), who hand-coded them using neurophysiological data available at the time, to control a lamprey-type agent. A similar approach is taken by Zheng et al. (2004) to model leech swimming. Others within the ALife community have taken the idea further by also incorporating evolution to derive the network architectures. Ijspeert and Kodjabachian (1999) applies a developmental as well as an evolutionary process in deriving the network structure, using swimming speed and muscular contortion for the fitness evaluation and a set of production rules for the developmental process. More realistic models include those of Sfakiotakis and Tsakiris (2006) who were able to replicate some of the biological movement data observed for the *Anguilla anguilla* eel (although their model omitted evolutionary mechanism). In the model, the 'eel' would navigate by incorporating sensory input from the front part of the animat. Beauregard and Kennedy (2006) further developed a model of an undulatory lamprey that could essentially track the movement of and follow, an object. This latter work was motivated out of a need to develop more realistic swimming algorithms for the computer animation industry. Our own model is explained in the following section.

Physical model

The Animat Fig. 1(a) represents a segment of the animat constructed out of layers. For clarity, not all springs have been depicted (in reality, each block in the animat has a 'crane-like' structure of springs to prevent it from collapsing in on itself). Since the animat is three dimensional, it is possible for it to undulate in multiple directions and/or demonstrate other types of movement depending on the output neurons of the neural network model.

The equations controlling the springs apply Hooke's Law with dampening dynamics; see Table 1 for the physical parameters of our system. Given a spring with mass points p_1 and p_2 on either end, it is compressed by forcing p_1 towards p_2 and vise-versa. Using p_1 as an example, the force exerted upon it by the internal dynamics of the spring, is computed as follows:

$$\overrightarrow{F_{p_1}} = -r \cdot \overrightarrow{V_{p1}} + k \cdot d, \tag{1}$$

where r is a dampening factor, $\overrightarrow{V_{p1}}$ is the velocity of p_1, k is a spring constant defining spring torque and d is the displacement of the spring from resting length. A change in the mass point's velocity, $\overrightarrow{V_{p1}}$, is governed by a change in its acceleration, $\overrightarrow{A_{p1}}$,

$$\overrightarrow{A_{p1}}(t + \Delta t) = \overrightarrow{A_{p1}}(t) + \frac{\overrightarrow{F_{p1}} + \overrightarrow{F_{p_1}^E} + \overrightarrow{F_{p_1}^W}}{m_{p1}}, \tag{2}$$

$$\overrightarrow{V_{p1}}(t + \Delta t) = \overrightarrow{V_{p1}}(t) + \overrightarrow{A_{p1}}(t + \Delta t) \cdot dt, \tag{3}$$

where m_{p_1} is the mass of p_1 and dt is the time-step (0.05) used during the integration process (20 integration steps). Note that $\overrightarrow{F_{p_1}^E}$ is an external force applied to the mass point whenever the output neuron controlling its associated spring, becomes activated, and $\overrightarrow{F_{p_1}^W}$ represents the current environmental force yielded by the surrounding 'water'. Finally, the position of the mass point, and hence the length of the spring is updated as follows,

$$\overrightarrow{P_{p1}}(t + \Delta t) = \overrightarrow{P_{p1}}(t) + \overrightarrow{V_{p1}}(t + \Delta t) \cdot dt. \tag{4}$$

The above equations afford a fluid and life-like representation.

The Environment The agent's environmental niche is modelled on movement through water. To keep things simple, we rely solely on the animat's current velocity to derive the environmental water force. This is the approach taken by most researchers (e.g. Sfakiotakis and Tsakiris (2006)). For all block faces, the water force, $\overrightarrow{F_W}$, is iteratively computed and applied to each constituent mass point,

$$\overrightarrow{F^W} = -\frac{1}{2} \cdot \nu \cdot \delta \cdot \alpha \cdot \overrightarrow{V} \cdot \left(\vec{V}\right)^2 \tag{5}$$

On the RHS of the equation, the velocity parameter, $\vec{V}$, is squared to give an indication of 'speed'. This determines the magnitude of the velocity, together with the viscosity, υ, and the drag, δ , in determining the amount of environmental force that should be applied; α is simply the area of the animat block face ($6.4 \cdot 0.35/8$). The environmental parameters that we use are given in Table 1.

Parameter	Value
Mass point masses	20.0
Layer springs	k=200,r=10.5
Block springs 'struts'	base k=25, r=50
Block springs 'crane'	k=500,r=50
Environmental viscosity, ν	10
Environmental drag, δ	1.0
Animat block count	8
Animat length	6.4
Animat width	0.35
Neurons per block	10

Table 1: Physical Parameters. Note that k=spring constant and r=spring dampenner. Note that the spring constants of the block 'struts' are controlled by the CTRNN. The base value of 25 sets an upper bound for one of these constants. All values are reflective of trial and error.

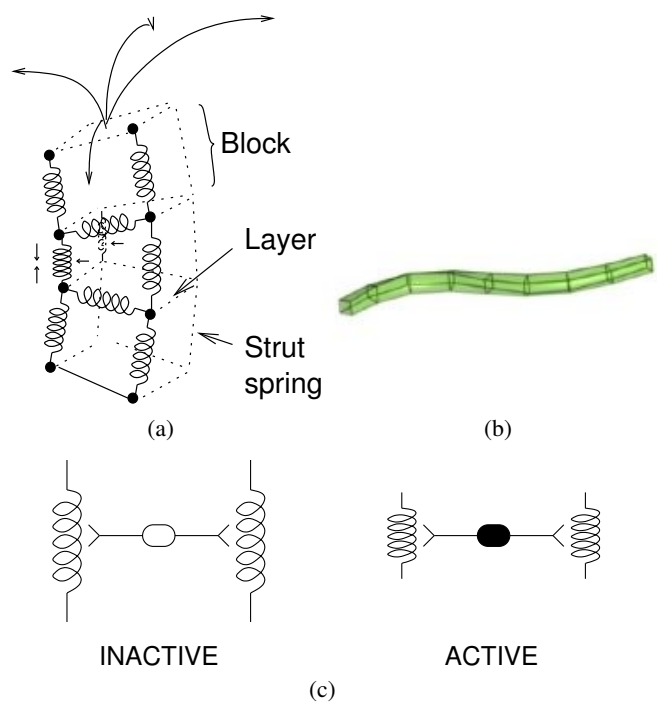

Figure 1: (a,c) Diagrams indicating how springs contract in pairs as activated by motor neurons; (b), a rendered visualization of the agent.

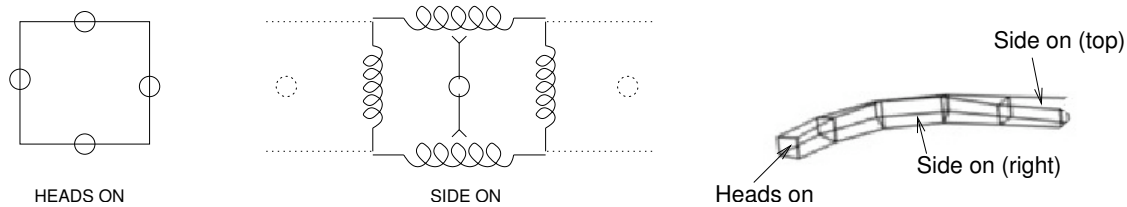

Figure 2: A diagram showing heads–on and side–on views of where the motor neurons (circles) are structurally located.

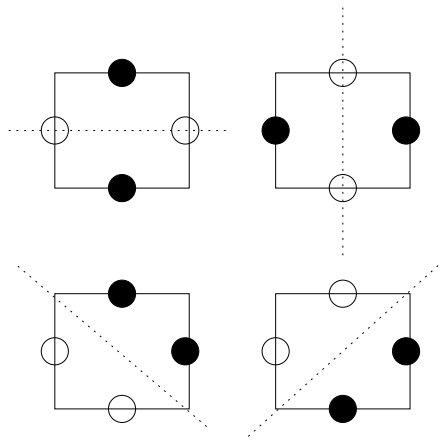

Figure 3: As additional gene values in our genome, different active motor configurations (those motors that play a part in movement) can be selected for during a process of evolution. Filled circle - active motor; dashed line - plane of symmetry. Taken from a heads–on perspective, looking down the animat from one of its ends.

Neural network implementation

A continuous time recurrent neural network (CTRNN) is employed to regulate the spring-pair compressions. The activation of a 'motor' neuron is used to calculate the spring constant of an associated spring, with a preset force of '200' but this force is only applied if the activation is between zero and one. The maximum spring constant is '25'. Each block is self-contained and houses the neural network architecture encoded for by the 'neural architecture' parameters labeled in Fig. 5. We could have chosen instead to encode a set of neural architecture parameters for each block, but this would have drastically increased the size of the search space during a process of evolution, reducing its tractability.

Weight values and inter-connectivity amongst neurons are entirely governed by neuron position. Neurons change position during evolution (because of mutation) except for the motor neurons which always reside within the centers of the block faces, see Fig. 2. Note further that the motor neurons can either bring about movement activity, or they can just serve as general interneurons. Accordingly, different 'active motor configurations' (Fig. 3) will have different impacts on the range of possible movements so are evolved along with the neural network architecture (see section 'Evolutionary Algorithm').

The membrane potential of a CTRNN neuron is computed according to its incoming pre-synaptic activity. In discrete time-steps, this activity, μ_i, of neuron i can be modelled as follows (based on Blynel and Floreano (2002)),

$$\mu_i(n+1) = \mu_i(n) + \frac{\left(-\mu_i(n) + \sum_{j=1}^{N} w_{ij} A_j(n) + I\right)}{\tau_i}, \tag{6}$$

where n is a discrete time-step and τ_i is the time constant for neuron i. The value A_j is the current output activity of presynaptic neuron j. The value I represents an external input current. Since the network never receives any 'sensory input', it has to be triggered and so we set this value to 1.0 for the first two neurons for the very first time-step of a simulation run. A neuron might also be inhibitory in which case the signs of all outgoing weights are flipped.

A weight value from neuron i to neuron j is derived according to the Euclidean distance between them, the impact of which is controlled by a parameter ξ (=2.0), Eq. 7. We also constrain the weights to fall within w^{max} (=20) and w^{min} (=0.0001), Eq. 8.

$$\lambda_{ij} = \frac{\xi}{d_{ij}}, \tag{7}$$

$$w_{ij} = \begin{cases} w^{max} & \lambda_{ij} \geq w^{max}, \\ w^{min} & \lambda_{ij} \leq w^{min}, \\ \lambda_{ij} & \text{otherwise.} \end{cases} \tag{8}$$

Connectivity Connectivity between a pair of neurons is established according to a minimum distance requirement. Hence we employ three threshold parameters. The first decides interneuron–interneuron connectivity; the second decides interneuron–effector connectivity and the third decides connectivity between neurons from a pair of contiguous subneural architectures. Since there are no sensory neurons currently employed in the model, there are no additional parameters as might be expected for more advanced architectures. A connection is formally decided as follows:

$$C_{ij} = \begin{cases} 1 & d_{ij} \leq \Gamma_q, \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

where Γ_q is a threshold parameter that we evolve, initialized to [0.01,2.5]. Note that in terms of connectivity between subnetwork architectures s_p and s_q, neuron *i* from s_p is only allowed to make one inter-subnetwork connection and that is explicitly chosen to be neuron *i* from s_q, see Fig. 4. Motor neurons never make inter-subnetwork connections.

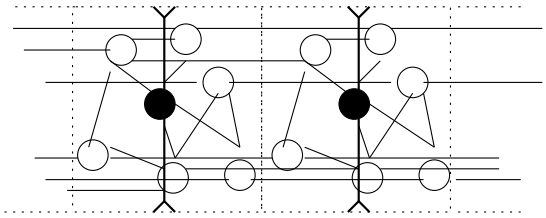

Figure 4: A diagram clarifying the repeated neuron architectures and how they are contiguously connected. Filled circles – motor neurons; unfilled circles – general interneurons. Non-dashed lines – interneuron connections.

Evolutionary Algorithm

The evolutionary algorithm optimizes the architectural parameters of the CTRNN network as described above together with the active motor configurations. These parameters form the individual's *genotype*. Both real and binary parameters are employed throughout so the algorithm employs a mixed real-valued and binary representation, see Fig. 5.

The implementation method that we employ utilizes self-adaptation of the mutation parameters. This affords a broader discovery of solutions during early evolution and a finer traversal during the later stages (e.g. Liang et al. (1998)).

Fitness measure. This is simply chosen to be the distance that the animat can move forwards during 200 time-steps. During testing of the simulation, we occasionally found that the physical interactions of the springs would oscillate out of control due to poor dampening dynamics causing either an animat 'implosion' or 'explosion'. Whenever this happened, the fitness of the individual would be set to -10000.

Mutation. All real-valued genes are mutated with values drawn from a normal distribution having an expectancy of 0 and a variance governed by the mutation parameter. This occurs for every gene with a preset probability, Φ, set to 0.02; when it occurs, the mutation parameter is also adapted. For a real-valued y-positional gene,

$$y_i = \begin{cases} y_i + N(0, \sigma_i) & \text{rand}() < \Phi, \\ y_i & \text{otherwise.} \end{cases} \tag{10}$$

whilst for a binary valued inhibitory or motor activity gene,

$$m_i = \begin{cases} !m_i & \text{rand}() < \Phi, \\ m_i & \text{otherwise.} \end{cases} \tag{11}$$

The adaptation of the mutation parameters relies on the setting of two strategy parameters, $\tau_o = 1.0/\sqrt{2D}$ and

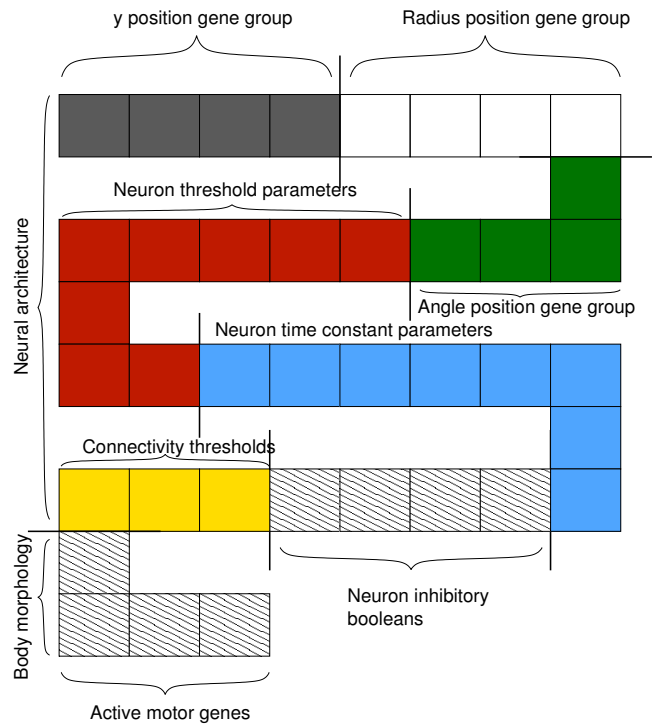

Figure 5: A representation of an individual chromosome where gene groups have been partitioned. The example is for an individual with 8 neurons per animat block. Note that in this example, there are only 4 genes per positional group because the positions of the four motor neurons remain fixed.

$\tau_1 = 1.0/\sqrt{2\sqrt{D}}$ which have been shown to be optimal in a process of self-adaptation (see Bäck and Schwefel (1993)). *D* is the dimensionality of the gene vector. Therefore, with respect to the example given in Fig. 5, *D*=4, for any of the positional groups; D=8 for the thresholds and time constants and lastly D=3 for the connectivity thresholds. The σ mutation parameters are then 'self-adapted' as shown,

$$\sigma_i \longleftarrow \sigma_i * \exp\left(N(0, \tau_0) + N_i(0, \tau_1)\right). \tag{12}$$

Crossover. All genes within a chromosome are subject to being exchanged with genes from another chromosome (*single point crossover*). This process occurs with a preset probability, χ, set to 0.2; when it occurs the mutation parameters are also crossed over between the same two chromosomes. Note that candidates for this operation are pulled out from the population at random, up to the size of the population. For any gene (both real-valued and binary types) the crossing over process can be summarized as follows:

$$\langle y_i, y_j \rangle = \begin{cases} \langle y_j, y_i \rangle & \text{rand}() < \chi, \\ \langle y_i, y_j \rangle & \text{otherwise.} \end{cases} \tag{13}$$

Selection. In our scheme we use binary tournament selection with an elitist strategy. To begin with, we rank the

population (of size 100) according to fitness and pick an elite number of individuals (=6) to form the start of the offspring. The remaining offspring population is then chosen randomly using binary tournament selection in which (until the offspring population reaches the population size), two population members are picked at random and the fittest is chosen with a preset probability (=0.9). Except for the elitist, all members are then subjected to the above mutation and crossing over operations. We use binary tournament selection since it facilitates diversity.

Results

Fig. 7 shows the progression of best fitness for a simulation run. Given the active motor configurations annotated on to the plot, we can see that fitter individuals favor a bilaterally symmetric configuration (motors opposite each other). This optimal configuration of up/down or left/right active motor configurations was found to emerge in six out of six simulation runs (results omitted). These configurations evolved with the neural network architecture as shown in Fig. 8. The neural network dynamics of the active motor neurons are given in Fig. 9. Interestingly, we can see that whilst all blocks demonstrate a variety of mostly CPG dynamics, only the first, second and sixth show both active neurons to have CPG dynamics. Blocks three, four, five, seven and eight all only show one of the active motor neurons to have CPG dynamics (either A or B); the other neuron for one of these blocks has an activity that is shown to trail off towards a negative value. Furthermore, this active motor neuron is seen to alternate in successive blocks for which only one of the neurons is active: block three's active motor neuron is motor neuron 'A', whilst block four's is 'B' and then block five's is 'A' again; block seven's is 'A' whilst block eight's is 'B'. These dynamics directly contribute to the movements of the animat since we know that the spring pairs for a given block compress whenever the associated motor neuron has an activity of between 0 and 1.

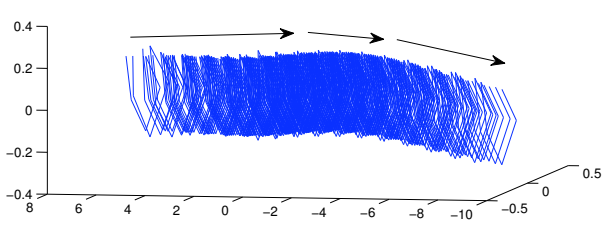

Figure 6: A motion–captured visualization of the animat travelling in the direction marked by arrows. Note, a more negative value on the lower axis indicates further forward travel.

Fig. 6 shows a motion–captured visualization of the animat travelling in the direction marked by arrows. It is difficult to observe any undulatory movements. In fact, the animat moves forward by crumpling and then extending its body segments (although there are also some very small undulatory-type movements).

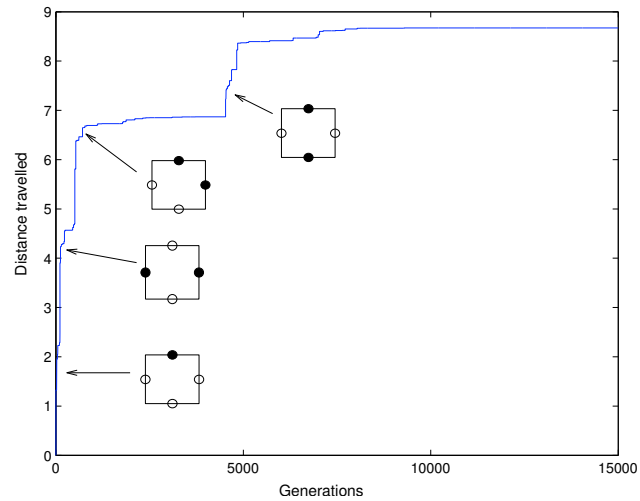

Figure 7: A graph showing the progression of the fittest population member over a simulated evolutionary period. Active motor configurations annotate different points of innovation; active motors are represented by filled circles.

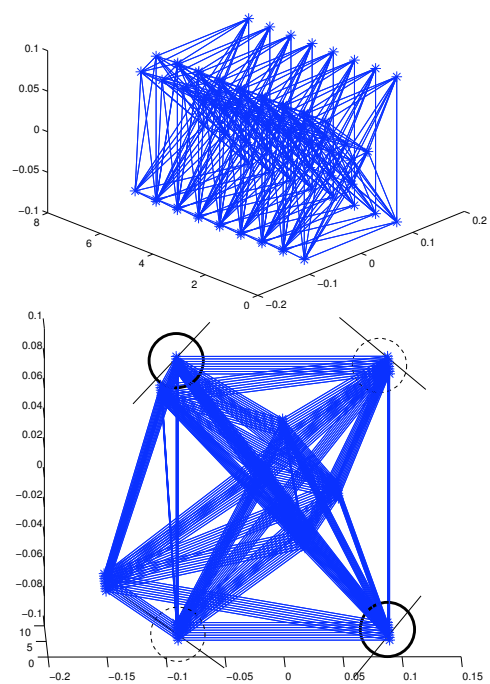

Figure 8: Visualizations of the neural network architecture from the fittest individual. In the lower visualization, the animat is shown head–on and the motor neurons that are used for movement are solidly circled (filled circles in Fig. 7); (motor neurons not used for movement – dashed circles).

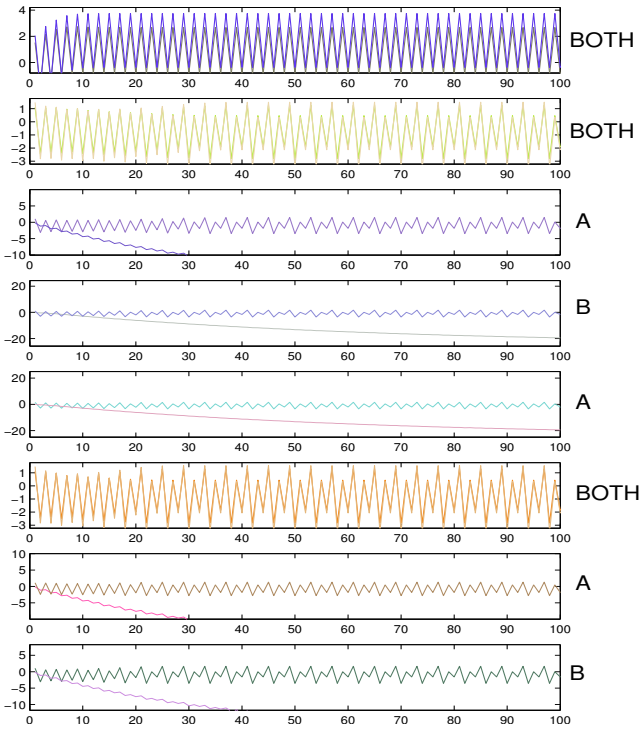

Figure 9: The CTRNN dynamics for the active motor neurons in each animat block, 1-8, each represented by a sub graph. They are labelled 'BOTH', 'A' or 'B' according to whether both active motor neurons or only one of them (A or B) demonstrate CPG dynamics.

Discussion

As noted in the introduction, nature has provided examples showing that nervous system architecture is coupled to body-plan morphology. Based on this knowledge and the hypothesis that different couplings are favored by different environments, we constructed a simple model to shed some light on how such a coupling could emerge during a process of simulated evolution.

The discovery that our simulated agent would in many ways reflect natural counterparts in terms of preferring a bilaterally symmetric motor configuration is interesting, since other than the physical characteristics of the agent (how the springs were interconnected) and the physical features of the environment (drag and viscosity), we placed no further constraints upon the movement mechanisms. Of course we intuitively know that bilateral symmetry is advantageous (consider how we walk), but the way that the control system should arrange itself – to *configure* itself – in concert with the body-plan, is less clear. The coupling is complex and the two components should not be considered separate; the architecture of the nervous system places indirect pressure on the type of body-plan morphology (configuration of motors, within our model), that can evolve and vise-versa. Our framework has helped us to elucidate this interplay of body, nervous system and environment. Our long term research goal is to better our understanding of this, especially with a fuller regard to a change in body-plan symmetry; with this paper, we have only begun to broach the subject of this major evolutionary transition.

We share the intuitive view that a change in body-plan symmetry likely occurred as organisms found themselves immersed in environments requiring directional movement. It is further our own view that this change would have been facilitated by an evolutionary drive towards those body-plan/nervous system couplings that minimize energy loss. Indeed, evolution could in part be pressured by those movement mechanisms requiring no energy. As an analogy, consider how we would compress a spring. One must of course apply energy, but upon releasing the compression, the spring passively returns to its natural resting length. Crucially, this 'self-stabilizing' process is inherent in muscle contractions and relaxations, (e.g. Pfeifer and Bongard (2006)).

In terms of our model, we can consider how simulated evolution strives to find a balance between the number of active spring compressions and the number of passive spring relaxations so that the above self-stabilizing process can be 'optimized'. If we attach an energy measure to this process, we might find that evolution prefers a maximization of passive relaxations, since this would perhaps conserve the most energy (but the springs would always have to be compressed first). However, defining such a measure will be hard, since in reality, energy can be lost from both the nervous system and from the 'muscles' and determining the levels of loss from each will directly determine the evolutionary process. Ideally, both energy losses will be coupled.

Future work

We are currently extending the model to address more fully the evolutionary transition from radial to bilateral symmetry. This will allow us to extensibly investigate the complex interactions between body, nervous system and environment and will bring us a step closer in answering why certain of these interactions emerge in a particular way. We plan to do this by (i) extending the range of morphological features; (ii) incorporating a more flexible body-plan/nervous system coupling representation; (iii) extending the flexibility of the environment so that at various stages of the simulation, specific couplings are pre-disposed.

Conclusion

In setting out to model an elongated agent that could move (e.g. swim) through water, we have shown an evolutionary preference for a bilaterally symmetric control system (a CTRNN) whose dynamics ultimately shape this movement mechanism. We further conclude that since the CTRNN architecture is coupled to the body-plan motor system, and that

movement depends on this coupling, forward movement requires a very specific coupling in order that the correct dynamics can be obtained; and moreover, evolution prefers a coupling that will, because of its inherent features, endow bilaterally symmetric functionality.

Acknowledgement

The first author is grateful to Honda Research Institute Europe for a PhD studentship that supports this work.

References

Bäck, T. and Schwefel, H.-P. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, 1(1):1–23.

Beauregard, M. and Kennedy, P. J. (2006). Robust simulation of lamprey tracking. In *Parallel Problem Solving from Nature - PPSN IX*, pages 641–650, Berlin. Springer-Verlag.

Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behaviour. *Adaptive Behaviour*, 1(1):91–122.

Blynel, J. and Floreano, D. (2002). Levels of dynamics and adaptive behavior in evolutionary neural controllers. In *From Animals to Animats 7: Proceedings of the seventh international conference on simulation of adaptive behavior*. MIT Press.

Bongard, J. C. and Paul, C. (2000). Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. In *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*. MIT Press.

Collins, A. G. and Valentine, J. W. (2001). Defining phyla: evolutionary pathways to metazoan body plans. *Evolution and Development*, 3(6):432–442.

Downing, K. L. (2007). Supplementing evolutionary developmental systems with abstract models of neurogenesis. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 990–996. ACM.

Eggenberger, P. (1997). Evolving morphologies of simulated 3d organisms based on differential gene expression. In *Proceedings of the Fourth European Conference on Artificial Life*, pages 205–213. MIT Press.

Ekeberg, O. (1993a). An Integrated Neuronal and Mechanical Model of Fish Swimming. In *Computation of Neurons and Neural Systems (CNS)*, pages 217–222. Klumer.

Ekeberg, O. (1993b). A combined neuronal and mechanical model of fish swimming. *Biological Cybernetics*, 69:363–374.

Gillis, G. B. (1996). Undulatory locomotion in elongated aquatic vertebrates: Angulliform swimming since Sir James Gray. *American Zoology*, 36:656–665.

Holland, N. D. (2003). Early central nervous system evolution: and era of skin brains? *Nature reviews*, 4.

Husbands, P., Smith, T., and O'Shea, M. (1998). Better living through chemistry: Evolving gasnets for robot control. *Connection Science*, 10(3–4):185–210.

Ijspeert, A. and Kodjabachian, J. (1999). Evolution and development of a central pattern generator for the swimming of a lamprey. *Artificial Life*, 5(3):247–269.

Lacalli, T. (1996). Dorsoventral axis inversion: a phylogenetic perspective. *BioEssays*, 18(3).

Liang, K.-H., Yao, X., Liu, Y., Newton, C. S., and Hoffman, D. (1998). An experimental investigation of self-adaptation in evolutionary programming. In *Lecture Notes in Computer Science*, volume 1447, pages 291–300. Springer-Verlag.

Meinhardt, H. (2002). The radial-symmetric hydra and the evolution of the bilateral body plan: an old body become a young brain. *BioEssays*, 24(2002):185–191.

Pfeifer, R. and Bongard, J. C. (2006). *How the Body Shapes the Way We Think, A New View of Intelligence*. MIT Press.

Psujek, S., Ames, J., and Beer, R. D. (2006). Connection and coordination: The interplay between architecture and dynamics in evolved model pattern generators. *Neural Computation*, 18:729–747.

Sfakiotakis, M. and Tsakiris, D. (2006). Simuun: A simulation environment for undulatory locomotion. *International Journal of Modelling and Simulation*.

Sims, K. (1994a). Evolving 3d morphology and behaviour by competition. In *ALife IV Proceedings*, pages 28–39. MIT Press.

Sims, K. (1994b). Evolving virtual creatures. In *Computer Graphics, Annual Conference Series*, pages 15–22.

Taylor, T. and Massey, C. (2001). Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life*, 7(1).

Zheng, M., Iwasaki, T., and Friesen, W. (2004). Systems approach to modeling the neuronal cpg for leech swimming. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 703–706.