Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology

accepted for publication in Adaptive Behavior

Auke Jan Ijspeert*, John Hallam*, David Willshaw**

*Dept. of Artificial Intelligence, U. of Edinburgh, 5 Forrest Hill, Edinburgh EH1 2QL

**Centre for Cognitive Science, U. of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW

{aukei, john}@dai.ed.ac.uk, david@cns.ed.ac.uk

Abstract

This paper presents how neural swimming controllers for a simulated lamprey can be developed using evolutionary algorithms. A genetic algorithm is used for evolving the architecture of a connectionist model which determines the muscular activity of a simulated body in interaction with water. This work is inspired by the biological model developed by Ekeberg which reproduces the central pattern generator observed in the real lamprey [Ekeberg 93]. In evolving artificial controllers, we demonstrate that a genetic algorithm can be an interesting design technique for neural controllers and that there exist alternative solutions to the biological connectivity. A variety of neural controllers are evolved which can produce the pattern of oscillations necessary for swimming. These patterns can be modulated through the external excitation applied to the network in order to vary the speed and the direction of swimming. The best evolved controllers cover larger ranges of frequencies, phase lags and speeds of swimming than Ekeberg's model. We also show that the same techniques for evolving artificial solutions can be interesting tools for developing neurobiological models. In particular, biologically plausible controllers can be developed with ranges of oscillation frequency much closer to those observed in the real lamprey than Ekeberg's hand-crafted model.

Keywords: Neural control; genetic algorithm; simulation; central pattern generator; swimming; lamprey.

1 Introduction

As locomotion is one of the most important aspects of living systems, the control of movements has been extensively studied both in biological sciences and engineering. Neurobiology has shown the diversity of effective motion controllers present in the animal kingdom. These controllers, composed of interconnected neurons, are able to create the oscillatory activity necessary for motion as well as to adapt to dynamic environments using sensory information. The controllers are generally distributed and present remarkable robustness and flexibility. Inspired by these interesting features, roboticists have applied artificial neural networks to control robots, using abstract mathematical models of neurons for processing the signals needed for control.

Advances in neurobiology and in robotics are now at a point where models of biological controllers can be directly implemented in robots [Beer *et al.* 93]; neural controllers inspired by insects have, for instance, been used to control the walking of hexapod robots [Espenschied *et al.* 93, Cruse *et al.* 95]. Robotics benefits from biology as a source of inspiration on how to solve control problems, and neurobiology may in return benefit from these experiments by receiving an evaluation of the capacity of a biological model to produce a specific behaviour.

Another interesting interaction is the use of artificial neural network techniques for investigating the computations performed by the central nervous system or for developing neurobiological models. An adapted backpropagation algorithm has, for instance, been used to develop a biologically plausible connectionist model reproducing the local bending reflex of the leech [Lockery & Sejnowski 93]. Evolutionary algorithms are another popular design technique for artificial neural networks [Schaffer *et al.* 92, Yao 93], which have recently been used for developing biophysical models of single cells [West & Wilcox 97].

Our research is interested in biologically inspired neural locomotion controllers for autonomous agents and how to develop controllers using a genetic algorithm (GA). In particular, we study here *anguiliform* swimming, one of the simplest forms of locomotion, which is used by eels and lampreys. Motion is generated by traveling undulations of the body, without the use of fins. Although simple, this type of locomotion presents the basic features of any natural motion control, namely the creation and the modulation of sets of oscillations with characteristic frequencies and phase lags.

This work is inspired by Ekeberg's connectionist model of the lamprey's swimming controller [Ekeberg 93]. Using a GA, we develop artificial controllers by evolving the connections and synaptic weights between neurons similar to those of Ekeberg's model. We are interested in evaluating how efficient GAs can be for designing neural networks which can produce the patterns of oscillation necessary for swimming, and whose activity can be modulated by simple input signals. A possible application of such a design technique could be, for instance, the automatic generation of neural controllers for a swimming robot. We are also interested in investigating whether there exist alternative neural configurations to the biological one found in the lamprey. Such a study may provide some insights into what are the necessary characteristics for central pattern generators in general. Finally, we investigate whether GAs could be used as a tool for neurobiological modeling, with the evolutionary process instantiating variables such as synaptic weights in a connectionist model whose connectivity and behaviour correspond to neurobiological observations.

In this paper, we firstly present the current knowledge of the lamprey's swimming controller and the connectionist model made by Ekeberg. We then develop alternative solutions in three evolutionary stages, with firstly the evolution of a segmental oscillator, secondly the evolution of the couplings between a set of segmental oscillators, and finally, the evolution of sensory feedback from stretch sensitive cells. As an example of how the GA could be used to develop biological models, we then repeat the evolutions by adding constraints on the encoding in order to evolve controllers which preserve the segmental connectivity observed in the lamprey.

2 Biological background and Ekeberg's connectionist model

The lamprey is one of the earliest and simplest vertebrates. It has no paired fins and swims by propagating an undulation along its body, from head to tail. The network

of neurons located in its spinal cord which controls the motion has been extensively studied [Buchanan & Grillner 87, Grillner *et al.* 88, Grillner *et al.* 91, Grillner *et al.* 95]. Physiological experiments on isolated spinal cords have shown that this network of neurons is a central pattern generator (CPG), which means that it can generate patterns of oscillation without oscillatory excitation either from the brain or from sensory feedback [Delcomyn 80, Grillner *et al.* 88, Getting 88]. Because small parts of the spinal cord can be isolated (up to 2 segments) and still produce oscillations when subjected to an excitatory bath, the CPG is considered to be an interconnection of local (segmental) oscillators. The complete controller is formed of approximately 100 of these segments. The CPG can produce oscillations at frequencies varying between 0.25 and 10 Hz; and it maintains constant wavelength of the undulation (usually corresponding to the length of the body) as the frequency is changed.[1]

The CPG controlling the swimming has been modelled at several levels of abstraction, namely at a biophysical, a connectionist and an abstract oscillator level. Biophysical models using relatively realistic neural models have demonstrated that the current state of knowledge of the lamprey is sufficient to produce models whose results agree with the physiological observations [Ekeberg *et al.* 91, Wallén *et al.* 92]. Connectionist models of the CPG [Buchanan 92, Williams 92, Ekeberg 93] use less realistic neuron models which capture only the main feature of neurons, that is, their ability to change the frequency of action potential spikes in their axon (the output) depending on the sum of activity in their dendrites (the total input). These have demonstrated that the connectivity in itself is sufficient to produce many of the oscillation patterns observed in the real lamprey without the need for complicated neuronal mechanisms. Finally, at the most abstract level, the swimming controller can be modeled as a chain of mathematical oscillators in order to study which kind of couplings can produce a phase relation between segments which is constant over the whole spinal cord of the lamprey and also remains a fixed proportion of a cycle when the frequency of the oscillations is changed [Kopell 95, Williams & Sigvardt 95].

The work presented here is situated at the connectionist level. A connectionist model has the advantages of being abstract enough to be tractable, in the sense that the simulation of the complete CPG (segments over the whole spinal cord) and of a mechanical simulation of the body of a lamprey do not require too much computation time and too many parameter choices. At the same time it is concrete enough to give some interesting insights into the functioning of the CPG of the real lamprey (as mentioned above, connectionist models can produce most of the oscillation patterns observed in the real lamprey). Finally, the connectionist models made for modelling the lamprey's CPG are very similar to those currently used for the control of some robots [Lewis *et al.* 93]. Therefore, learning how to design them may be directly applicable to the design of a controller for a swimming robot, for instance.

We have reproduced Ekeberg's two-dimensional connectionist model [Ekeberg 93]. This model has inspired our work and it is the biological model with which our evolved controllers will be compared. We will refer to it as the *biological model*. The model combines a simulation of the neural controller with a mechanical simulation of a lamprey in water allowing direct evaluation of how neural activity is transformed into mechanical movements. It also offers the possibility of investigating the effect of sensory feedback

[1] This maintenance of constant wavelength over a large range of frequencies is an interesting aspect of the CPG. It means that the phase lag between segments is maintained at a constant proportion of the period of oscillations and that the time delays between segments vary considerably depending on the frequency. The lags between segments can not therefore be due to delays in conduction along descending axons because axons have, more or less, fixed velocities of conduction.

from stretch sensitive cells. We present here the model and give some quantitative results (Ekeberg's paper [*ibid.*] presents mainly qualitative behaviours of the model) which will be compared with the performances of the evolved swimming controllers.

The neural controller is composed of 100 interconnected segmental networks (Figure 1). Each segmental network is an oscillator made of two motoneurons (MN), two excitatory interneurons (EIN), two contralateral inhibitory interneurons (CIN) and two lateral inhibitory interneurons (LIN). Each neuron unit in the model represents a population of functionally similar neurons in the real lamprey. Neurons receive excitatory input from the brainstem. The weights of the segmental connections are given in Table 1. The complete controller is formed by interconnecting 100 copies of this segmental network. An interconnection consists of extending the connection from one neuron to another in one segment to the corresponding neuron in neighbouring segments (Table 1). Because the interconnections in the lamprey are not well known, Ekeberg has simplified them by choosing a symmetric extension (in the caudal and rostral direction) for all connections except for the extensions from CIN which project more caudally.

A neuron unit is modeled as a leaky integrator with a saturating transfer function. Its output u corresponds to the mean firing frequency of the population it represents ($\in [0,1]$) and is calculated as follows:

$$\dot{\xi_+} = \frac{1}{\tau_D}\Big(\sum_{i\in\Psi_+} u_i w_i - \xi_+\Big) \tag{1}$$

$$\dot{\xi_-} = \frac{1}{\tau_D}\Big(\sum_{i\in\Psi_-} u_i w_i - \xi_-\Big) \tag{2}$$

$$\dot{\vartheta} = \frac{1}{\tau_A}(u - \vartheta) \tag{3}$$

$$u = \begin{cases} 1 - \exp\{(\Theta - \xi_+)\Gamma\} - \xi_- - \mu\vartheta & (u > 0) \\ 0 & (u \leq 0) \end{cases} \tag{4}$$

where w_i are the synaptic weights, Ψ_+ and Ψ_- represent the groups of pre-synaptic excitatory and inhibitory neurons respectively, ξ_+ and ξ_- are the delayed 'reactions' to excitatory and inhibitory input and ϑ represents the frequency adaptation observed in some real neurons. The parameters of each type of neuron are given in Table 2. These parameters, as well as the connection weights of Table 1, have been defined so that the simulation of the model fits physiological observations [Ekeberg 93].

When the segmental networks receive adequate excitation from the brainstem, they oscillate regularly with the left and right neurons out of phase. The frequency of oscillations increases with the level of excitation. The interconnections between segments leads to traveling waves of oscillations when the most rostral segments receive extra excitation compared to the others. The lag between segments relative to the period of oscillation is then constant over the whole spinal cord, except for the first and last ten segments at the extremities where it decreases towards zero. The value of the lag varies with the amount of extra excitation, and the higher the extra excitation, the larger the lag (or the shorter the wavelength). Interestingly, the frequency of oscillation and the wavelength of the undulation can be changed nearly independently. The model can therefore reproduce the capacity of the real lamprey to cover a whole range of different frequencies of oscillations while keeping the wavelength of the undulation constant. Our simulations show that frequencies between 1.6 and 5.5 [Hz] and phase lags per segment between 0.0% (no travelling wave) and 2.4% of the period can be obtained (the wavelength is then smaller than 50% of the 100-segment body).

We have also reproduced Ekeberg's mechanical model of the lamprey's body interacting with water [Ekeberg 93]. The body is made of 10 rigid links connected through one-degree of freedom joints, with each mechanical link corresponding therefore to 10 neural segments. Muscles are connected to each link and are modeled as a combination of springs and dampers. The motoneurons can 'contract' muscles by increasing their spring constant, thus reducing their resting length. Propulsion through water is obtained when the traveling waves of neural activity are transformed into traveling undulations of the body. The head to tail propagation of the undulation leads to forces from the surrounding water which propel the fish forward. We have simulated the biological model over the whole range of frequencies and phase lags the neural controller can produce, and the highest speed the simulated lamprey can reach is 0.50 m/s (with oscillations at 5.5 Hz and relative lags between segments of 1.2%).[2]

The mechanical simulation allows a direct evaluation of the efficiency of the CPG to control swimming (in terms of speed of swimming, for instance). It also allows study of the effect of sensory feedback from the edge cells, which are stretch-sensitive cells located on both sides of the spinal cord [Viana Di Prisco *et al.* 90]. These cells can be modeled as outputing a signal proportional to the local curvature of the body. While in [Ekeberg 93] the effects of the edge cells on the controller are negligible (very low synaptic weight), in [Ekeberg *et al.* 95], it is shown that feedback can help the lamprey to cross a speed barrier (water with local speed opposite to the direction of swimming) by coordinating the neural activity with the actual movements of the body and preventing a change of direction.

3 Evolving central pattern generators

We develop alternative swimming controllers by using the GA to define suitable connections and synaptic weights between neurons similar to those of Ekeberg's model. The similarity of the neuron models will allow us to directly compare the performances of the evolved solutions with those of the biological model. The controllers are evolved in three stages. First, segmental oscillators are evolved; then multi-segmental controllers are generated by evolving the couplings between copies of a chosen segmental oscillator; and, finally, connections providing sensory feedback from stretch-sensitive cells are added. The first stage requires only the simulation of the neural activity within a segment, while the two last are realised with a simulation of the whole controller together with the mechanical simulation of the body interacting with water. This decomposition into stages was motivated by properties of the biological controller we wanted to reproduce. Firstly, segmental networks can oscillate independently when isolated from the others and from sensory feedback. Secondly, being a CPG, the multi-segmental controller can produce coordinated traveling waves of oscillations in a completely isolated spinal cord indicating that there is no need for sensory feedback for the intersegmental coordination. Finally, it appears that the role of feedback from stretch sensitive cells is mainly necessary for coordinating the neural activity with the actual movements of the body when these are disturbed by the environment. The third stage therefore includes the simulation of a speed barrier which can only be crossed when sensory feedback is provided to the CPG.

[2]We have corrected some of Ekeberg's mechanical values and therefore obtain speeds of swimming which are approximately 40% smaller for the two quantitative examples given in [Ekeberg 93]. However, the qualitative behaviours of the two simulations are identical. See [Ijspeert *et al.* 98a] for more details.

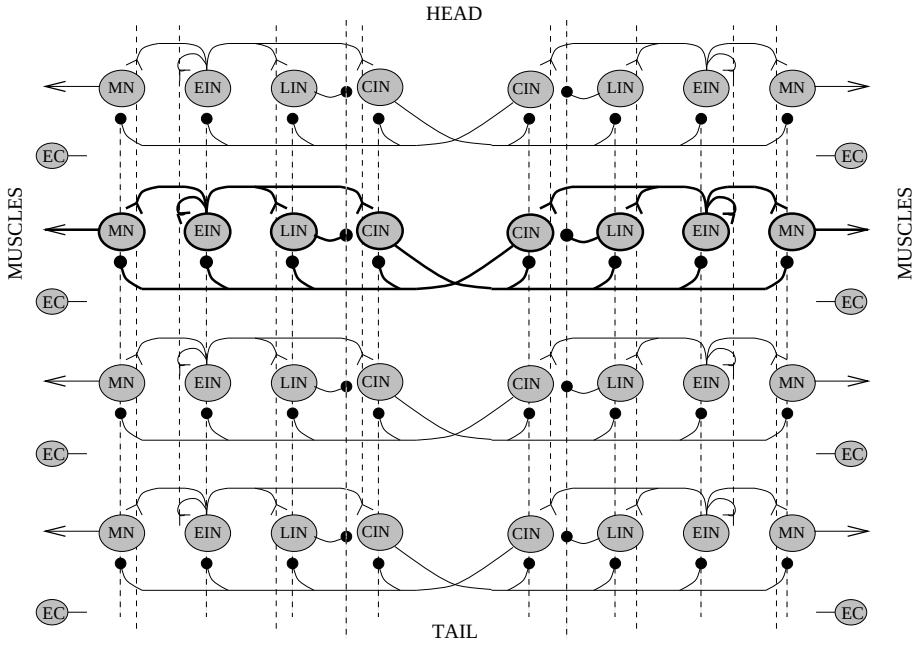

Figure 1: Biological controller. The controller is made of 100 interconnected segmental oscillators (only 4 segments shown) composed of 8 neurons each. Four types of neurons are present in the oscillators: 3 types of interneurons EIN, CIN and LIN and the motoneurons MN. The controller can receive feedback from the stretch sensitive edge cells EC. The dashed lines indicate the projections from segmental connections to neigbouring segments. See text and Table 1 for more explanations.

from: / *to:*	EINl	CINl	LINl	EINr	CINr	LINr	BS
EINl	0.4 [2, 2]	-	-	-	−2.0 [1, 10]	-	2.0
CINl	3.0 [2, 2]	-	−1.0 [5, 5]	-	−2.0 [1, 10]	-	7.0
LINl	13.0 [5, 5]	-	-	-	−1.0 [1, 10]	-	5.0
MNl	1.0 [5, 5]	-	-	-	−2.0 [5, 5]	-	5.0
EINr	-	−2.0 [1, 10]	-	0.4 [2, 2]	-	-	2.0
CINr	-	−2.0 [1, 10]	-	3.0 [2, 2]	-	−1.0 [5, 5]	7.0
LINr	-	−1.0 [1, 10]	-	13.0 [5, 5]	-	-	5.0
MNr	-	−2.0 [5, 5]	-	1.0 [5, 5]	-	-	5.0

Table 1: Biological configuration, as given in [Ekeberg 93]. Excitatory and inhibitory connections are represented by positive and negative weights respectively. Left and right neurons are indicated by l and r. *BS* stands for brain stem. The number of segments to which connections from a particular segment extend, in the rostral and caudal directions respectively, are given in brackets.

Neuron type	Θ	Γ	τ_D	μ	τ_A
EIN	−0.2	1.8	30 ms	0.3	400 ms
CIN	0.5	1.0	20 ms	0.3	200 ms
LIN	8.0	0.5	50 ms	0.0	-
MN	0.1	0.3	20 ms	0.0	-

Table 2: Neuron parameters, as given in [Ekeberg 93]. Θ is the threshold, Γ the gain, τ_D the time constant of the dendritic sums, μ the coefficient of frequency adaptation and τ_A, the time constant of the frequency adaptation.

The controllers are evaluated depending on their ability to control swimming. In order to obtain similar performance to the biological controller, we define the following goals for the evolved controllers:

- production of stable oscillations over the spinal cord,
- generation of phase lags between segments,
- frequency of oscillation dependent on global excitation,
- lag between segments dependent on the proportion of extra excitation of the most rostral segments,
- independence of lag and frequency,
- large ranges of frequencies, lags and speeds of swimming,
- capacity to incorporate sensory feedback from stretch sensitive cells for crossing a speed barrier.

3.1 First evolutionary stage: Segmental oscillators

A segmental oscillator is generated by evolving both the connectivity between the 8 neurons composing the biological segment and the *sign* (excitatory and inhibitory) of each interneuron. The quality of a solution is evaluated depending on its ability to produce regular oscillations whose frequency can be varied with the level of excitation.

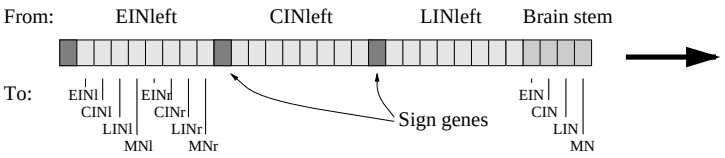

	EINl	CINl	LINl	EINr	CINr	LINr	Brain Stem
EINl	1.3	-0.8			-2.7		4.3
CINl		-3.5	-4.2				10.8
LINl	7.4						3.0
MNl	6.0				-3.0		5.8
EINr				1.3	-0.8		4.3
CINr		-2.7			-3.5	-4.2	10.8
LINr				7.4			3.0
MNr		-3.0		6.0			5.8

Figure 2: Encoding of a segmental network.

Encoding A chromosome represents the connectivity between the 8 neurons of the biological model. The number of neurons and their dynamics (defined by the 5 parameters of equations 1-4) are thus fixed and only the sign of neurons and the synaptic weights of the connections are evolved. We impose a left-right symmetry. A chromosome is a string of 31 genes which are real numbers between 0 and 1 (Figure 2). There are three *sign genes* which define whether the interneurons are inhibitory or excitatory.[3] The motoneurons are excitatory. The other genes directly encode the synaptic weights and correspond (via a linear transformation) to a real value between -5 and 0 or between 0 and 15, depending on the value of the sign gene (the interneuron is inhibitory if the sign gene is smaller than

[3]We chose to have neurons which are either inhibitory or excitatory (rather than neurons which could both inhibit and excite other neurons) in order to develop networks under the same constraints as the biological model, in the sense that biological neurons do not usually emit both excitatory and inhibitory neurotransmitters.

0.5 and excitatory otherwise). Finally the four last genes of a chromosome determine the synaptic weights of the connections coming from the brain stem and correspond to a real value between -5 and 15. The weight boundaries (-5 and 15) have been defined so that they include the range of weights of the biological model (-2 and 13). Note that, although we will keep the denominations EIN, CIN and LIN to distinguish between the different types of interneurons (the different neuron dynamics given by equations 1-4), they lose their meaning as a description of the function or even the sign these neurons had in the biological model.

Genetic Algorithm The algorithm used is a variation of the standard GA (see for instance [Goldberg 89]) with the usual binary encoding being replaced by a real number encoding. The algorithm starts with a randomly generated initial population. At each generation, a fixed number of children are generated using three operators, *crossover*, *mutation* and *pruning*. The *crossover* operator chooses couples of parents, with a rank-based probability, for breeding and creates couples of children either by two-point crossover (probability *Prob_Xover*) or by simply copying the two parents. The *mutation* operator mutates each gene of the children with a probability *Prob_mut* and the mutation consists of adding or substracting a small random number within a mutation range:

$$new\_value = old\_value + Mut\_Range \cdot rand$$

where *rand* is a random number $\in [-0.5, 0.5]$. Finally, the *pruning* operator is applied which randomly prunes a connection (probability *Prob_prun*) by setting the gene to the value corresponding to a null weight. The children are then evaluated, and the size of the population is kept constant by rejecting, at each generation, the worst solutions of the increased population (old population plus children). The GA parameters used for the evolution of segmental oscillators are given in Table 3.

Population size	100
Number of children	30
Crossover probability	0.5
Mutation probability	0.4
Mutation range	0.2
Pruning probability	0.1

Table 3: GA parameters for evolving segmental oscillators

Evaluation We define a fitness function that rewards three desired characteristics for segmental oscillators:

1. The production of regular oscillations of the motoneurons, with one peak of activity per period and with the left and right neurons out of phase.

2. A frequency of oscillation which can be varied and which increases monotonically with the level of external excitation.

3. A minimal set of connections. In [Ijspeert *et al.* 97], we evolved fully connected solutions, the weakest connections of which proved unnecessary for the creation of oscillations and could be removed without affecting the neural activity.

The mathematical definition of the fitness function is given in Appendix A1. Solutions are evaluated by fixed-duration simulations (simulated time of 3000 ms) with asymmetric initial conditions (all left neurons excited). If a network oscillates regularly, simulations with different levels of external excitation are carried out in order to determine its range of frequencies. Note that the evaluation of the fitness of a segmental oscillator is only based on the neural activity of the motoneurons because they are the only neurons influencing the muscular activity. This means that the interneurons can have any activity and that solutions with fewer than 8 active neurons can be developed.

Results We carried out 10 evolutions with populations of 100 chromosomes. The evolutions were stopped after 500 generations, an arbitrarily chosen limit which was sufficient for the best solutions of each run to reach a fitness value higher than 0.3. Each evolution started with a different randomly generated initial population. When the neurons are randomly connected, there is little chance of the resulting network oscillating. In our case, only 3 solutions within the 1000 initially generated solutions produce regular oscillations, and this only at a fixed frequency.[4] Within 500 generations, all evolutions converged to final populations composed of networks producing regular oscillations with variable frequency. Similary to the biological model, the frequency of oscillations and the amplitude of the motoneuron signals increase with the level of tonic excitation. When the oscillator receives too little excitation, the motoneurons stay inactive, and when the excitation level is too high they stabilize, after a few oscillations, in an asymmetrical state with one active and one inactive motoneuron. The fittest solutions of all runs cover a larger range of frequencies than the biological model, reaching lower and higher frequencies. The results are summarized in Table 4.

The ten runs did not all converge to the same network configuration, but rather to several different solutions with similar fitness values (see [Ijspeert *et al.* 98a] for a detailed description of all the evolved configurations). These best solutions differ not only in their connectivity, but also in the sign of the interneurons and in the number of neurons active in the creation of oscillations. Six evolutions converged to networks with only 6 rather than 8 active neurons. These solutions are composed of the two motoneurons and four inhibitory interneurons of two different types which create the oscillations. This demonstrates that excitatory interneurons are not necessary (at least in a connectionist model) for the production of oscillations over a large frequency range. Most often the inactive neurons are the LINs, which is probably due to their high threshold; they have more chance to be inactive unless they receive some strong excitatory input. Amongst these six runs, four (runs 3,6,8 and 10) converged to similar configurations based on an identical oscillator structure (when the possible permutations due to the left-right symmetry are taken into account) composed of inhibitory EIN and CIN and the motoneurons (Figure 3). Although there are some variations of the values of the connection weights, the solutions have all the same connections with similar strengths and present the same behaviour. Figure 4 shows the simulation and the configuration of one of them.

Although most of the evolved oscillators differ significantly from the biological segmental network, one controller (the best solution of run4) presents strong similarities with it. That solution is composed of the same interneurons — same number and with the same sign — as the biological model. Furthermore, although it has more connections

[4]Note that the random configurations which do not oscillate still receive different fitness values depending on how much the amplitudes of their motoneurons vary before reaching the equilibrium state.

	Fitness value	Frequency range in [Hz]	Oscillating neurons	Number of connections
Biol. model	0.18	[1.7, 5.6]	EIN+,CIN−,LIN−	26
run1	0.44	[1.0, 8.4]	EIN−,CIN+,LIN−	32
run2	0.39	[1.0, 6.8]	EIN−,LIN−	34 → 24
run3	0.43	[1.1, 7.7]	EIN−,CIN−	34 → 26
run4	0.47	[1.2, 9.9]	EIN+,CIN−,LIN−	38
run5	0.40	[1.1, 7.9]	EIN−,CIN−,LIN−	34
run6	0.51	[1.1, 11.3]	EIN−,CIN−	40 → 24
run7	0.54	[1.2, 8.0]	EIN−,CIN−	32 → 24
run8	0.31	[1.1, 5.8]	EIN−,CIN−	36 → 22
run9	0.60	[1.0, 10.4]	EIN−,CIN+,LIN−	32
run10	0.44	[1.3, 8.6]	EIN−,CIN−	28 → 22

Table 4: Summary of results for the evolved segmental networks. The table gives, for the best solution of each population, the lowest and the highest frequency at which the segmental network can oscillate, the interneurons active in the oscillator with the sign of their influence and the number of connections of the oscillator. When fewer than 6 interneurons are active, the number of connections after complete removal of the inactive neurons are given.

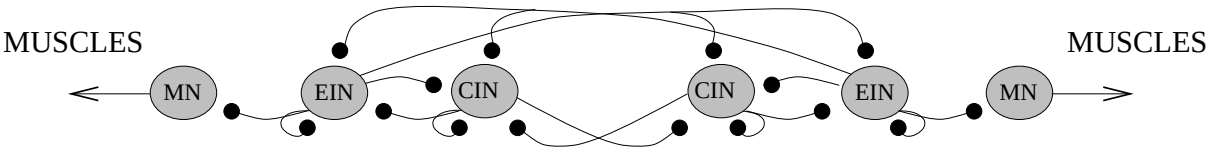

Figure 3: Segmental network configuration to which four out of ten runs converged (runs 3,6,8 and 10).

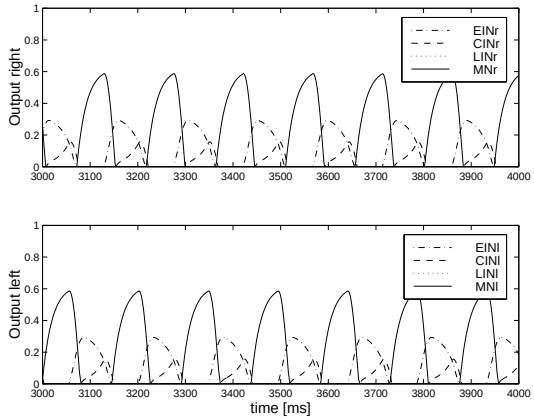

from: / *to:*	EINl	CINl	LINl	EINr	CINr	LINr	BS
EINl	−1.9	−5.0	-	−4.2	−0.1	-	12.1
CCINl	−2.8	−1.7	-	−5.0	−2.1	-	4.7
LINl	-	-	-	-	-	-	-
MNl	−5.0	-	-	-	-	-	4.8
EINr	−4.2	−0.1	-	−1.9	−5.0	-	12.1
CCINr	−5.0	−2.1	-	−2.8	−1.7	-	4.7
LINr	-	-	-	-	-	-	-
MNl	-	-	-	−5.0	-	-	4.8

Figure 4: *Left:* oscillations in an evolved segmental network (run 6). *Right:* corresponding neural configuration.

than the biological model, it has a similar structure with inhibitory connections from the CIN to all the neurons on the contralateral side and a strong inhibitory connection from the LIN neurons to the CIN neurons on the same side acting as a burst terminator. As with the biological model, the LIN becomes active later in the cycle and inhibits the ispilateral CIN allowing the activity on the other side to start.

The pruning operator of the GA has led to solutions using significantly fewer than the 56 possible connections. The solutions have therefore few weak (low weight) connections which do not play an important role in the creation of oscillations. In the case of solutions which use fewer than 8 neurons, the majority of the connections to and from the inactive neurons have been cut; for these solutions Table 4 gives the number of connections before and after the removal of the inactive neurons.[5]

3.2 Second evolutionary stage: Multi-segmental controllers

We generate complete CPGs by evolving the interconnections between fixed segmental networks. The best segmental oscillators of each evolution of the previous evolutionary stage are chosen. The evaluation of complete CPGs is based on simulations of both the neural activity and the mechanical movements of the body. The quality of the solutions will depend on their capacity to control swimming at different speeds, frequencies of oscillation, and lags between segments.

Encoding and GA A chromosome encodes how the segments are interconnected. Similarly to the biological model, interconnections are extensions of segmental connections to the corresponding post-synaptic neurons in neighbouring segments. For each segmental connection, the extensions in each direction are encoded into a chromosome with values between 0 (no extension to neighbour segments) and a maximal extension of 12 (this value has been chosen to include the maximal extension of the biological model which is 10). The number of genes is therefore twice the number of segmental connections. As previously, a left-right symmetry is assumed. The connection weights of the segmental networks are fixed except for a rescaling depending on the extension of the interconnection: the weight of a connection to a neuron is divided by the number of segments it receives input from.

The same GA as for the evolution of segmental oscillators is used but without the pruning operator. Although the genes are transformed into integers when the chromosome is decoded, they are real numbers in the GA. The GA parameters for the evolutions are given in Table 5.

Population size	40
Number of children	12
Crossover probability	0.5
Mutation probability	0.4
Mutation range	0.2

Table 5: GA parameters for evolving multi-segmental controllers.

[5]The fact that there are still some weak connections which could be removed without affecting the oscillation is due to the fact that the mutation operator can at any time give a non-zero value to a weight which had previously been pruned (set to zero). The mutation operator therefore slows down the convergence to the minimal set of connections.

Evaluation Multi-segmental controllers are evaluated for their capacity to control swimming in the mechanical simulation. The required features for the controllers are:

1. generation of stable oscillations in the 100 segments with coordinated phase differences for the creation of traveling undulations of the body,

2. ability to change the speed of swimming by changing either the frequency of oscillation or the wavelength of the undulation,

3. ability to change the frequency and the wavelength independently, by changing, respectively, the global excitation level and the amount of extra excitation on the most rostral segments.

The mathematical definition of the fitness function is given in Appendix A2. The evaluation of a controller consists of several simulations with different levels of global excitation and extra excitation of the most rostral segments in order to determine their influence on the frequency, the phase lag between segments and the speed. In order to allow comparisons with the biological controller, a special emphasis is given to the capacity to swim with a wavelength corresponding to the length of the body (phase lag per segment of 1%), and therefore the capacity to change the frequency over a large range is only measured for lags close to 1%.

Results We carried out 10 evolutions based on the best evolved segmental oscillators of the previous evolutionary stage. Each population is made of 40 chromosomes and evolutions are stopped when the fitness of the fittest solution of the population ceases to increase significantly (between 75 and 300 generations depending on the run).

The random coupling of the initial populations has a different effect depending on the segmental network. Some segmental networks are quite robust and produce stable oscillations despite the random couplings, but most of them fail to produce stable oscillations. A general observation is that, among the randomly connected segmental oscillators which can produce stable oscillations, very few can have the phase between segments varied with the extra excitation on the first segments.

All runs successfully converged to controllers capable of generating stable oscillations with phase lags between segments for relatively large ranges of frequency and phase lags. Similarly to the biological model, the evolved CPGs produce lags between segments which are constant over the whole spinal cord, except for the extremities (because of the lack of extensions from the boundaries, the segments at the extremities tend to oscillate in phase). The lags increase when extra excitation is applied to the most rostral segments. After the evolutions, we tested the best solutions of each run over a whole range of levels of excitation and extra excitation in order to determine their exact ranges of speed, frequency and phase lag. The results are summarized in Table 6.

The multisegmental controllers cover approximately the same range of frequency as the corresponding isolated segmental networks. Five evolved controllers cover a larger range of phase lags than Ekeberg's solution, with a maximum lag of 3.8% for the solution of run4. The speed range covered by the evolved solutions is also generally larger than the range of Ekeberg's model (8 solutions reach higher speeds). Note that some solutions swim slowly backwards when all segments oscillate in phase (no extra excitation on the first segments), because of the kind of wriggling they then perform. Also, some solutions produce lags even without extra excitation of the first segments.

	Fitness value	Frequency range in [Hz]	Relative lag range in [%]	Speed range in [m/s]
Biol. model	0.08	[1.6, 5.6]	[0.0, 2.4]	[−0.03, 0.50]
run1	0.07	[1.9, 8.5]	[0.0, 1.9]	[−0.04, 0.59]
run2	0.09	[1.0, 6.8]	[0.1, 3.1]	[0.06, 0.58]
run3	0.07	[1.1, 7.8]	[0.0, 2.0]	[0.08, 0.49]
run4	0.13	[0.8, 10.1]	[0.0, 3.8]	[−0.03, 0.53]
run5	0.26	[1.1, 8.4]	[0.0, 3.6]	[−0.10, 0.55]
run6	0.02	[1.1, 7.0]	[0.0, 1.3]	[0.07, 0.51]
run7	0.02	[1.2, 7.5]	[0.0, 2.0]	[−0.03, 0.52]
run8	0.10	[1.1, 5.8]	[0.0, 3.0]	[−0.03, 0.56]
run9	0.13	[2.0, 11.0]	[−1.7, 3.5]	[−0.07, 0.50]
run10	0.18	[3.3, 8.7]	[0.3, 2.2]	[0.18, 0.60]

Table 6: Summary of results for the evolved multi-segmental controllers. The table gives, for the best solution of each population, the ranges of frequency, phase lag and speed it can produce.

Although most solutions cover larger ranges of frequencies and phase lags than the biological model, the independence of control of these variables is usually less good. For several solutions, for instance, increasing the global level of excitation not only increases the frequency of oscillations but also influences the relative lag between segments (usually decreasing it). This problem is illustrated by the best solution of run9 which can reach high frequencies and high lags, but not both together, which explains why its maximum speed is relatively low (a combination of high frequency and high lag is necessary for high speeds in our mechanical simulation).[6] Our fitness function has not prevented this relative lack of independence because, although it rewards independence, it does it only for the small area of the global excitation versus extra excitation space that is visited for the evaluation. The dependence occurs mainly at high levels of global excitation and extra excitation which are not visited for the evaluation. This relative problem could probably be solved by sampling more combinations of both types of excitation for the evaluation.

There are no clear conclusions to make from the way the segmental networks are coupled (see [Ijspeert *et al.* 98a] for a detailed description of all the evolved configurations). The extents of the interconnections vary considerably from one solution to another, and there is, for instance, no systematic coupling favouring one direction. The extents of the coupling and the asymmetries vary depending on the types of the connected neurons, the influence of the connection, the phase differences between the activities of the neurons within the segmental oscillator, *etc.* A few general observations can, however, be made. On average, the ipsilateral connections between neurons of the same type (the self-connections) extend more rostrally than caudally while the contralateral connections between neurons of the same type extend more caudally. Also a large majority of the inhibitory connections sent to the motoneurons extend more caudally. Finally the CPGs based on similar segmental oscillators have also similar couplings, with the same asymmetries of coupling for most segmental connections.

Figure 5 shows a simulation of one of the evolved multi-segmental controllers. Interestingly, the evolved CPGs can also perform turning although this aspect has not been taken into account in the evaluation functions. Similarly to the biological controller, giving a higher excitation level to one side of the spinal cord leads to a difference of the motoneuron

[6]Note that the best solution of run9 also has the idiosyncrasy of producing, for some levels of global excitation, a caudal to rostral traveling wave when it receives no extra excitation. As soon as some extra excitation is given, the undulation travels from head to tail.

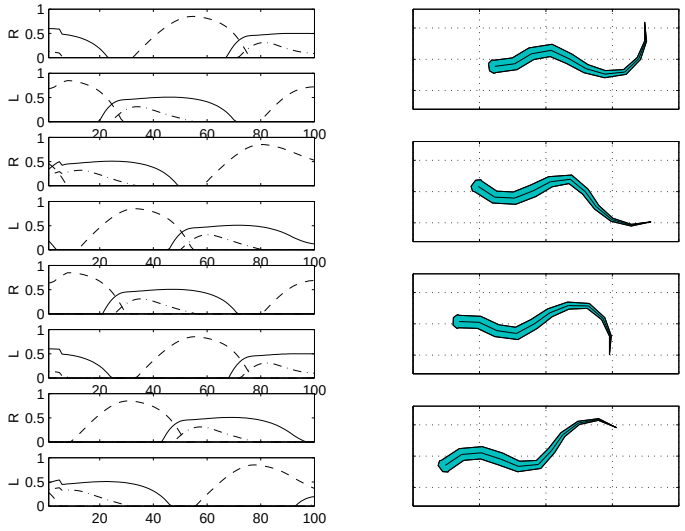

Figure 5: Simulation of one of the evolved swimming controllers (best of run7). The frequency of oscillations is 4.0 Hz, the phase lag between segments is 1.1% of a cycle duration and the speed of swimming is 0.41 m/s. *Left:* Neural simulation, the continuous line correspond to the MN, the dashed line corresponds to the CIN and dotted-dashed line to the EIN. *Right:* Mechanical simulation, there are 60 ms between each snapshot, and the dotted vertical lines are separated by 100 mm.

amplitude and, because of stronger muscle contractions on one side, to a change in the direction of swimming. In the case of the evolved controllers, some left-right permutations of the interneurons may be necessary for optimising the turning ability (because of the left-right symmetry, these permutations have no effect on straight swimming). For instance, for avoiding antagonistic effects on the motoneuron amplitude when extra excitation is applied, the neurons sending inhibitory connections to a motoneuron should be on the contralateral side. Finally, initial tests on some of the evolved controllers show that it is possible to produce backwards swimming by giving more excitation to the most caudal segments and therefore inducing a tail to head undulation, but only very limited lags and speeds can then be obtained.

3.3 Third evolutionary stage: Sensory feedback from stretch sensitive cells

The CPGs evolved so far work in an open loop, in the sense that they receive no feedback from the mechanical simulation. In the lamprey, sensory feedback is provided to the CPG by stretch sensitive cells (the edge cells) situated on both side of the spinal cord. There are both inhibitory and excitatory edge cells and they emit signals proportional to their elongation [Viana Di Prisco *et al.* 90]. These cells are useful for coordinating the neural activity with the actual body movements. Ekeberg showed in particular that this sensory feedback is necessary for crossing a barrier of water with local speed opposite to the direction of swimming [Ekeberg *et al.* 95]. We show here a preliminary experiment on how to evolve this sensory feedback for the artificial controllers in order to allow them to cross a similar speed barrier.

Encoding and GA A chromosome encodes how each segmental oscillator receives sensory feedback from two stretch sensitive cells situated at the sides of the corresponding part of the spinal cord. It encodes the synaptic weights of the connections from one edge cell to the 8 (or 6) neurons of the corresponding segmental network. We impose symmetry between both sides of the segmental network as well as over the whole spinal cord (the 100 segments). Each connection can be either excitatory or inhibitory and genes are decoded into synaptic weights between -2 and 5. Note that we have fixed these boundaries to be relatively small in order to prevent the sensory feedback disturbing the production of oscillations. The same GA with the same parameters as for the first evolutionary stage is used (Table 3), but without the pruning operator.

Evaluation Solutions are rewarded depending on their ability to cross a speed barrier. The barrier is 150 mm wide (half the length of the simulated lamprey) and the speed of the water flow opposite to direction of swimming is set to be 40% higher than the speed of swimming of the lamprey. The fitness function is defined to reward:

1. progression through the barrier in the direction of swimming,
2. minimal deviation in the direction,
3. minimal difference of speed between swimming with and without sensory feedback.

This last point was added to prevent important changes in the swimming patterns when sensory feedback is given to the controller. Evaluations are realised at a chosen combination of frequency of oscillations and phase lag between the segments (approximately in the middle of their respective range). The mathematical definition of the fitness function is given in Appendix A3.

Results We carried out 5 runs (populations of 40 chromosomes) using the best controller of run7 of the previous evolutionary stage. This controller was chosen because it combines relative large ranges of frequencies, lags and speeds and an independence of control of the frequency and the lag similar to the biological model. Sensory feedback is evolved to all 6 neurons forming the segmental oscillator (2 EINs, 2 CINs and 2 MNs). The chosen levels of excitation produce swimming (without sensory feedback) at a frequency of oscillations of 4.0Hz, a relative lag of 1.1% and a corresponding speed of 0.41 m/s. The speed of the barrier was set to 0.57 m/s (40% higher than the speed of swimming). Without sensory feedback, the lamprey is not able to cross the speed barrier (Figure 6, top). The increase in local forces at the entrance to the barrier disturbs the undulation of the body and leads to extra bending of the upper part of the body which forces the lamprey to change its direction of swimming. The lamprey eventually swims perpendicularly to its initial direction.

All runs converged within 80 generations to controllers able to cross the speed barrier (Figure 6, bottom). The connections from the stretch sensitive cells are not identical in all the evolved controllers. All solutions have, however, evolved excitatory connections from the stretch sensitive cells to the interneurons inhibiting the motoneurons on the contralateral side, therefore preventing excessive bending. The solutions have also developed strong excitatory connections from the stretch sensitive cells to the motoneurons. Interestingly, these connections are not only ipsilateral (to prevent excessive bending) but also contralateral. The overall effect is that the amplitudes of the motoneuron signals are approximately 20% larger than without sensory feedback. The excitatory feedback to the

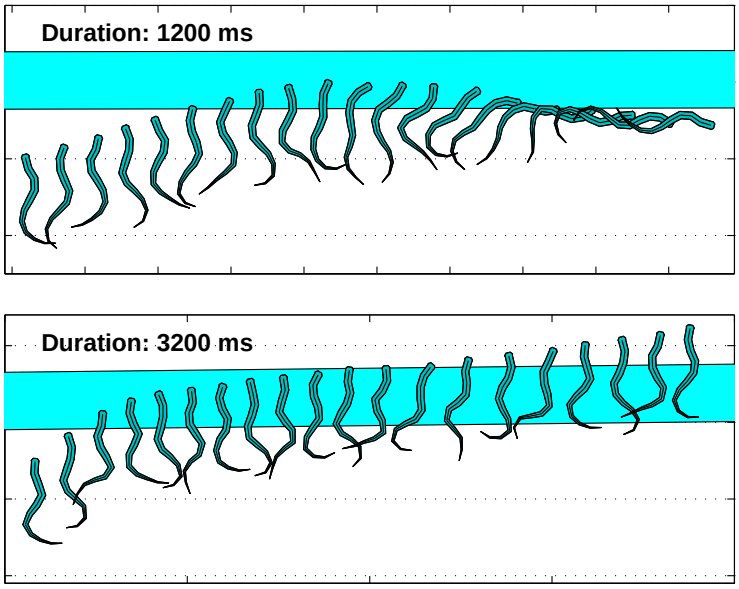

Figure 6: Effect of a speed barrier on swimming without (*top*) and with (*bottom*) evolved sensory feedback.

motoneurons leads thus to a stiffer body of the lamprey, which is therefore less sensitive to the extra forces due to the barrier. Note that it is possible (see [Ijspeert *et al.* 98a]) to evolve controllers which can cross a speed barrier without relying on an increase of the motoneuron amplitudes by evolving sensory feedback only to the interneurons. The crossing is then more difficult, however, and we had to reduce the speed of the barrier to obtain successful controllers (30% of the swimming speed). In that case, the sensory feedback acts only as a burst terminator which tends to switch the neural activity from one side to the other when the latter is over extended.

3.4 GAs as tool for neurobiological modeling

The method presented previously can also be used to develop connectionist models which present more similarities with the circuitry found in the lamprey. We believe that a GA could be a useful tool for neurobiologists to develop models when it is used to automatically generate instantiations of variables which are difficult to measure. As an illustration of this, we repeated the evolutions, with the same evolutionary stages and the same fitness functions, and restricted the space of possible solutions to controllers which preserve the segmental connectivity and the sign of neurons observed in the real lamprey. We here summarize the results of those evolutions; for more information, see [Ijspeert *et al.* 98a, Ijspeert *et al.* 98b].

We firstly evolved segmental oscillators by evolving the synaptic weights of a network with a fixed configuration corresponding to the biological segmental connectivity. We found that Ekeberg's set of synaptic weights could be modified in order to optimize the frequency range. Frequency ranges much closer to those observed in real lampreys (between 0.25 and 10 Hz) could be obtained (the best solution covered a range between 0.9 and 11.0 Hz). The main observation is that the increase of the frequency range is obtained through an increase of the strength of some inhibitory connections (the contralateral CIN to LIN and ispilateral LIN to CIN connections).

In the second evolutionary stage, we evolved the couplings between segmental oscillators corresponding to those of Ekeberg's biological model. As these coupling are not perfectly known in the real lamprey, the idea was to study which kind of interconnections can produce phase lags between segments which are constant over the spinal cord and independent of the frequency of oscillation. There are opposing views on the origin of these phase lags in the real lamprey, in particular whether they are due to differences in the intrinsic frequencies of the oscillators or to the nature of the coupling, and, in the latter case, whether the dominant coupling is ascending or descending (see for instance [Matsushima & Grillner 92, Wadden *et al.* 97]). We evolved interconnections with the same fitness function as for the artificial controllers and developed couplings which could produce the kind of phase lags observed in the real lamprey. The evolved solutions cover large ranges of lags and can reach slightly higher speeds than with Ekeberg's coupling. The evolved couplings in the five different runs we carried out (populations of 40 chromosomes) were all very similar, and significantly different from the simplified coupling of Ekeberg's model. We find that the evolved CPGs seem to have no dominant coupling (for instance, no systematic asymmetry) and they present both rostral and caudal asymmetries of projection depending on the segmental connection. As for Ekeberg's coupling, the CPGs rely both on the nature of the couplings and on differences of intrinsic frequencies for creating travelling waves. The coupling is optimized for waves to travel from head to tail and the extra excitation of the most rostral segments determines the exact phase lag (giving more excitation to those segments amounts to increasing their intrinsic frequency).

Finally, we also evolved sensory feedback to Ekeberg's controller which enabled the simulated lamprey to cross a speed barrier. Interestingly, the evolved connections correspond very closely to those observed in the real lamprey, with all connections from excitatory and inhibitory stretch sensitive cells established in the real lamprey [Viana Di Prisco *et al.* 90] being also present with the same sign in the evolved controllers. Crossing a speed barrier may therefore be a good example of the situation for which sensory feedback has been developed through natural evolution for the real lamprey.

4 Discussion

4.1 Discussion of the Method

Staged evolution Controllers were evolved in different stages in order to simplify the problem by decomposing it into subproblems. The decomposition was motivated by similar decompositions in the real lamprey and gave the possibility of direct comparison between evolved solutions and the biological model. One advantage of such a decomposition is to reduce the time needed to evolve a complete controller by avoiding evaluation of a whole controller made of segmental oscillators which do not oscillate correctly even when isolated. Decomposition presents the risk, however, of misleading the evolution as the fittest solutions of one stage may not necessarily be the best starting elements for the next stage. The segmental oscillators evolved in our first stage, for instance, have not been rewarded on their capacity to be interconnected, which may have constituted a handicap for the next stages.

Encoding scheme The simple encoding scheme which directly encodes synaptic weights and extents of interconnections into a chromosome meant that search was performed in a well defined and limited space of possible solutions (fixed neuron models, fixed maximum

number of neurons and left-right symmetry). Despite these limitations, the multidimensional search space was still large enough to generate a variety of interesting central pattern generators. One way to enlarge the search space could be to include the parameters of the neuron dynamics into the encoding. Initial tests show that the GA is able to evolve both these neuron parameters and the connectivity and generate working CPGs. Another interesting extension is to use an indirect encoding, such as a developmental program, which allows modularity and does not limit the number of neurons (see [Kodjabachian & Meyer 95] for a review). In [Ijspeert & Kodjabachian 98], we report how such a developmental encoding can be used to evolve swimming controllers for the lamprey in a single evolutionary stage.

Definition of the fitness functions Defining the fitness functions probably constitutes the most important part of the method. Our objectives were to develop CPGs capable of producing stable oscillations and of modulating the speed of swimming by varying, through a few external signals, the frequency and the phase lag of the oscillations. Once the objectives were defined, relatively little iteration was needed to define a suitable set of fitness functions. The functions are products of factors between 0 and 1 rewarding qualitative and quantitative aspects, which means that solutions which perform evenly in all aspects receive the highest fitness value. Both neural (e.g. the MN frequency) and mechanical (e.g. the speed of swimming) variables are rewarded, with the importance of neural aspects gradually decreasing and those of mechanical aspects increasing over the evolutionary stages. Ideally, only mechanical aspects should be examined as they alone effectively determine the quality of a swimming controller; rewarding neural activities presents the risk of biasing the search to particular network architectures. However, rewarding only mechanical variables would not have allowed the staged evolutionary approach discussed above.

GA and connectionist models GAs present several interesting features for the design of connectionist models compared to traditional learning algorithms. Learning algorithms for dynamical neural networks such as variations of the backpropagation algorithm, (see [Pearlmutter 95], for instance), iteratively update synaptic weights given an error function which calculates the difference between the current and the desired state trajectories of the network. These algorithms compute the gradient of the error function in terms of the synaptic weights and update the weights accordingly (gradient search). GAs are comparatively easier to implement and to use. Firstly, the fitness function does not need to be differentiable or even continuous. Secondly, there is no need to provide a specific oscillation (limit cycle) that the network should learn. Learning algorithms are restricted to problems in which the desired state trajectory is known in advance, which is not the case for many control problems. In our case, using a GA has allowed us to characterise the desired behaviour of the network at a "higher" level, for instance in terms of the range of frequencies which can be produced or in terms of the speed of swimming of the mechanical simulation. Also, the parallel and stochastic nature of GAs makes them less likely to be trapped in a local optimum than the gradient search of learning algorithms. Finally, GAs are particularly flexible and the type of evolved controllers can be modified at three levels: the encoding (e.g. by adding or removing constraints on the connectivity), the GA operators (e.g. adding a pruning operator) or the fitness function (e.g. adding or removing factors). These properties of GAs and evolutionary algorithms have made them a popular method for designing neural controllers for animats [Beer & Gallagher 92, Gruau 95, Kodjabachian & Meyer 98].

GA for neurobiological modeling The fact that the GA was able to find sets of synaptic weights and interconnections for CPGs when the search space was restricted to solutions presenting the biological segmental connectivity showed that it has the potential to be a useful tool for developing neurobiological models. The evolutionary process can indeed be used to generate a significant part (the synaptic weights and the intersegmental connections) of the connectionist model that Ekeberg set up by hand. Although it is possible, in this case, to define satisfactory values by hand, the GA proved useful for searching the different possibilities and for optimizing the synaptic weights in order to produce a larger range of frequencies. Several researchers in neurobiology are currently using evolutionary algorithms in a similar way, either for setting synaptic weights of neural network models [Eurich *et al.* 95, Eurich *et al.* 97], or for determining conductance parameters of biophysical models of single neurons [West & Wilcox 97]. From a general point of view, note that, although the GA can be very useful for demonstrating that a model can produce a specific behaviour (by finding efficient sets of unknown variables), it is less useful for invalidating a hypothetical model as an inability to find successful variable instantiations may be due to failings of the model or to problems with the GA set up, or both.

4.2 Discussion of the evolved controllers

Several interesting CPGs have been successfully evolved. A first observation is thus that there exists a variety of possible connectivities which produce the neural activities necessary for swimming.[7] We have demonstrated that there exist solutions alternative to the biological CPG which can control swimming with the same efficiency. These solutions differ from Ekeberg's biological model in terms of which neurons are interconnected, of the excitatory or inhibitory influence of the interneurons and even of the number of neurons involved in the rhythmogenesis.

Segmental oscillators We have evolved a diversity of artificial segmental oscillators with significantly larger ranges of frequencies than Ekeberg's segmental oscillator. Although one oscillator has a similar structure to the biological one, most evolved networks have interneurons that play other roles than in the biological network and have other signs. There is therefore no preferred role and sign for each dynamics of the different neuron models. A majority of oscillators are based on only four inhibitory interneurons, showing that, at least at the connectionist level, excitatory interneurons are not necessary for the production of oscillations over large frequency ranges.

Multi-segmental controllers The evolution of multi-segmental controllers showed that segmental networks could be interconnected in several ways and produce stable oscillations with the necessary phase lags for swimming. Similary to the biological model, swimming at different speeds and with different wavelengths of undulations can be produced by varying simple external signals (the excitation applied to the CPG). Turning can also be induced when one side of the controller receives more excitation than the other. The evolved controllers can generally cover larger ranges of frequency, phase and speed than Ekeberg's model but with less independence in the control of the frequency and the phase between segments. A general observation on the coupling between segments is that

[7] A similar observation was made by Lockery for the leech bending reflex, where many different networks could produce a physiologically accurate local bending input-output function [Lockery & Sejnowski 93].

there is no systematic asymmetry of interconnection favouring one direction but rather asymmetries in both directions which depend on the segmental connection. Note that the coupling of both Ekeberg's and the evolved controllers is a little bit crude in the sense that, except for the segments at the extremities, the synaptic weights of the extensions of a segmental connection are the same for all connected segments. It could be interesting to evolve, for each coupling, a normalising function which would determine the synaptic strengths depending on the length of the projection, in order to allow, for instance, a decrease of the influence of a segment on another with the distance separating them.

Sensory feedback The evolutions of connections from stretch sensitive cells illustrates how sensory feedback can be used by the evolved CPGs for coordinating the neural activity with the actual movements of the body and allowing the lamprey to cross a speed barrier. The feedback increases the stiffness of the body and acts as a burst terminator which switches the neural activity to the side which is excessively stretched. These preliminary experiments were realised with fixed external commands to the CPG (swimming straight) and the crossing of the barrier was therefore due to the passive properties of the CPG with sensory feedback. As active corrections in the commands may also help the lamprey to cross the barrier, further experiments should include these in order to determine the respective importance of sensory feedback and higher commands from the brain for crossing a barrier (this point was not considered in [Ekeberg *et al.* 95]).

Higher control centers As the speed and the direction of swimming can be modulated through two input signals, the movements of the simulated lamprey (in a plane) can be determined. The evolved controllers can therefore be used by higher neural centers for the generation of simple behaviours, with the two input nodes of the CPG providing the interface between the CPG and the higher control centers. It would therefore be interesting to extend the simulation by adding a simple simulated visual system to it and to evolve higher neural centers for simple sensory-motor coordination experiments such as avoiding obstacles or approaching and attaching to a host (the lamprey is a parasite).

5 Conclusion

This paper has described how a genetic algorithm can be used to develop connectionist neural networks for the control of anguiliform swimming. Inspired by Ekeberg's connectionist model of the central pattern generator of the lamprey, alternative artificial solutions have been successfully evolved. This work shows firstly that the genetic algorithm can be an interesting design technique for dynamical recurrent networks applied to locomotion control. With the technique, we have evolved networks which efficiently control the swimming of a simulated lamprey by producing stable oscillations whose frequencies and phase lags can be modulated by simple external signals. This work also illustrates that the biological connectivity is not unique and that it is possible to connect the interneurons of Ekeberg's model differently and obtain other efficient solutions. In particular, alternative solutions which cover larger frequency, lag and speed ranges can be obtained. Finally, we show how the same technique for evolving artificial solutions could potentially be used as a tool for developing neurobiological models. By incorporating biological knowledge of the real lamprey into the encoding and the evaluation function of the genetic algorithm, the evolutionary process can be used to find variables which are difficult to determine, in this case synaptic weights, for the production of a desired behaviour. Solutions which

satisfy the biological constraints and produce oscillations over a larger range of frequency than the original biological model have, for instance, been successfully evolved.

Acknowledgements: Auke Jan Ijspeert is supported by grants from the Swiss National Research Fund and from the European Commission (Marie Curie Fellowship). Facilities were provided by the University of Edinburgh. Many thanks to Aude Billard and Richard Reeve for comments on earlier versions of the manuscript.

A Fitness functions

A.1 Fitness of segmental oscillators

$$Fitness = \mathit{fit\_oscil} \cdot \mathit{fit\_connectivity} \cdot \mathit{fit\_freq} \in [(0.05)^5, 1]$$

where

$$\mathit{fit\_oscil} = \mathit{fit1} \cdot \mathit{fit2} \cdot \mathit{fit3} \in [(0.05)^3, 1]$$

$$\mathit{fit\_connectivity} = \begin{cases} -1.357 \cdot \mathit{con\_ratio} + 1.407 & \text{if } \mathit{con\_ratio} \in [0.3, 1] \\ 0.05 & \text{if } \mathit{con\_ratio} = 1 \\ 1 & \text{if } \mathit{con\_ratio} < 0.3 \end{cases}$$

$$\mathit{fit\_freq} = \begin{cases} 0.05 & \text{if } \mathit{fit\_oscil} < 0.45, \text{ otherwise:} \\ \begin{cases} 0.0864 \cdot \frac{\mathit{max\_freq}}{\mathit{min\_freq}} - 0.0364 & \text{if } \frac{\mathit{max\_freq}}{\mathit{min\_freq}} \in [1, 12] \\ 0.05 & \text{if } \frac{\mathit{max\_freq}}{\mathit{min\_freq}} = 1 \\ 1 & \text{if } \frac{\mathit{max\_freq}}{\mathit{min\_freq}} > 12 \end{cases} \end{cases}$$

The function *fit_oscil* is made of three factors which reward activity of the motoneurons which is varying (*fit1*), regular (*fit2*) and out of phase between left and right motoneurons (*fit3*). These factors are functions of statistical measures of the motoneuron signals (such as the standard deviation of the neuron output, the number of zeros, the difference of area under the output curves,...). They are bounded between 0.05 (bad behaviour) and 1 (good behaviour) and vary linearly between these values for the corresponding variables. The bounds for each variable have been determined by hand on 40 examples of different behaviours from initial experiments. It is possible to fix these bounds such that *fit_oscil* clearly makes the difference between interesting solutions and the others. A limit of 0.45 is thus determined above which a solution is certain to oscillate regularly with opposite behaviour between left and right motoneurons.[8] The function *fit_connectivity* rewards solutions with reduced connectivity. It is bounded between 0.05 and 1 and varies linearly between these values depending on the connectivity ratio (*con_ratio*), the ratio between the number of connections and the maximum number of possible connections (56). The smaller the connectivity the higher the reward, with a maximum reward of 1 for solutions with a connectivity ratio smaller than 0.3 (fewer than 18 connections). The factor *fit_freq* rewards solutions which can oscillate over a large range of frequencies when the excitation from the brain stem is varied. The range of frequencies is only measured if the solution oscillates regularly and with opposite behaviour between left and right (*fit_oscil* > 0.45) at the default level of excitation. The frequency range, the ratio between maximum and minimum frequencies, is measured by making a set of simulations at different levels of excitation. Only ranges in which the frequency and the amplitude of the motoneuron signals increase monotonically with the excitation level are rewarded.[9] Note that because

[8] Of the 40 examples of neural activities, 17 correspond to interesting behaviours. With the chosen fixed bounds of *fit1, fit2* and *fit3*, all the good behaviours have a *fit_oscil* value higher than 0.6 and all the others have a value lower than 0.25, with most lower than 0.1.

[9] The condition that the amplitude of the neural activity should not decrease when the frequency increases is added to prevent an antagonistic effect on the speed of swimming, as tests with the mechanical simulation showed that, in general, increases of the frequency or the amplitude of the motoneuron signals increase the speed of swimming.

the duration of a simulation is fixed and that a minimum of oscillations is required for verifying the stability of the signals, the lowest measurable frequency is approximately 1.0 Hz.

A.2 Fitness of multi-segmental controllers

$$Fitness = min\_fit\_oscil \cdot fit\_lagcontrol \cdot fit\_freqcontrol \cdot fit\_speed$$

where

- *min_fit_oscil* is the minimum *fit_oscil* value between segments 1,10,20,..,100 when the controller is simulated (simulations of 3000 ms) with an excitation level of *excit0*, corresponding to the middle of the range of excitations with which the segmental oscillator can oscillate, and without extra excitation on the first segments.
- $fit\_lagcontrol = \begin{cases} 0.05 + \frac{lag\_range1}{1+freq\_range1} & \text{if } < 1 \\ 1 & \text{otherwise} \end{cases}$
  Lag_range1 and *freq_range1* are measured by making several simulations at a fixed level of excitation (*excit0*) and with an increasing amount of extra excitation. The lag range is non-zero only if the oscillations are regular in all segments (*min_fit_oscil*>0.45) and if the lag increases monotonically with the amount of extra excitation. The lag range is divided by the frequency range in order to reward solutions which can change the lag between segments independently of the frequency.
- $fit\_freqcontrol = \begin{cases} 0.05 + \frac{freq\_range2}{1+lag\_range2} & \text{if } < 1 \\ 1 & \text{otherwise} \end{cases}$
  Freq_range2 and *lag_range2* are measured by making several simulations with excitation levels varying around *excit0* and with a fixed amount of extra excitation *extra0*. The value *extra0* is calculated from the previous lag measurement if *lag_range1* includes a lag of 1%. *Extra0* is then defined by taking the value of extra excitation corresponding to the lag closest to 1%. The frequency range is non-zero only if the value *extra0* exists, and if the frequency increases monotonically with the level of excitation.
- $fit\_speed = \begin{cases} 0.05 + speed\_range & \text{if } < 1 \\ 1 & \text{otherwise} \end{cases}$
  Speed_range corresponds to the range of speeds covered by all the simulations made for the definition of *lag_freqcontrol* and *fit_freqcontrol*.

The different value ranges are measured as the difference between the maximum and the minimum value and are normalised by a value corresponding to a target value. The targets have been fixed to 2.5% for the lag ranges, 10 Hz for the frequency ranges and 0.8 m/s for the speed range.

A.3 Fitness for sensory feedback

$$Fitness = fit\_progression \cdot fit\_deviation \cdot fit\_speed\_constancy$$

where

- *fit_progression* varies linearly between (and is limited to) 0.05 and 1 when the maximum progression of the head of the lamprey varies between 0 and 450 mm (width of the speed barrier plus the body length). The progression is measured as the distance (in the direction of swimming) between the head of the lamprey and the entry point to the barrier,

- *fit_deviation* varies linearly between 0.05 and 1 as the cosine of the deviation angle (measured between the initial direction of swimming and the direction after entry into the barrier) varies between 0 and 1,

- *fit_speed_constancy* varies linearly between 0.05 and 1 when the relative difference between the speeds with and without sensory feedback varies between 50% and 0%.

References

[Beer & Gallagher 92] R.D. Beer and J.C. Gallagher. Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.

[Beer *et al.* 93] R.D. Beer, R.E. Ritzmann, and T.M. McKenna. *Biological neural networks in invertebrate neuroethology and robotics.* Academic Press, 1993.

[Buchanan & Grillner 87] J.T. Buchanan and S. Grillner. Newly identified 'glutamate interneurons' and their role in locomotion in the lamprey spinal cord. *Science*, 236:312–314, 1987.

[Buchanan 92] J.T. Buchanan. Neural network simulations of coupled locomotor oscillators in the lamprey spinal cord. *Biological Cybernetics*, 66:367–374, 1992.

[Cruse *et al.* 95] H. Cruse, D.E. Brunn, Ch. Bartling, J. Dean, M. Dreifert, T. Kindermann, and J. Schmitz. Walking: A complex behavior controlled by simple networks. *Adaptive Behavior*, 3(4):385–418, 1995.

[Delcomyn 80] F. Delcomyn. Neural basis for rhytmic behaviour in animals. *Science*, 210:492–498, 1980.

[Ekeberg 93] Ö. Ekeberg. A combined neuronal and mechanical model of fish swimming. *Biological Cybernetics*, 69:363–374, 1993.

[Ekeberg *et al.* 91] Ö. Ekeberg, P. Wallén, L. Brodin, and A. Lansner. A computer-based model for realistic simulations of neural networks i: The single neuron and synaptic interaction. *Biological Cybernetics*, 65:81–90, 1991.

[Ekeberg *et al.* 95] Ö. Ekeberg, A. Lansner, and S. Grillner. The neural control of fish swimming studied through numerical simulations. *Adaptive Behavior*, 3(4):363–384, 1995.

[Espenschied *et al.* 93] K.S. Espenschied, R. Quinn, H.J. Chiel, and R.D. Beer. Leg coordination mechanisms in the stick insect applied to hexapod robot locomotion. *Adaptive Behavior*, 1(4):455–468, 1993.

[Eurich *et al.* 95] W.C. Eurich, G. Roth, H. Schwegler, and W. Wiggers. Simulander: a neural network model for the orientation movement of salamanders. *Journal of comparative physiology*, 176:379–389, 1995.

[Eurich *et al.* 97] W.C. Eurich, H. Schwegler, and R. Woesler. Coarse coding: applications to the visual system of salamanders. *Biological Cybernetics*, 77:41–47, 1997.

[Getting 88] P.A. Getting. Comparative analysis of invertebrate central pattern generators. In A. H. Cohen, S. Rossignol, and S. Grillner, editors, *Neural control of rhythmic movements in vertebrates.* Jon Wiley & Sons, 1988.

[Goldberg 89] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, 1989.

[Grillner *et al.* 88] S. Grillner, Buchanan J.T., P. Wallén, and L. Brodin. Neural control of locomotion in lower vertebrates. In A. H. Cohen, S. Rossignol, and S. Grillner, editors, *Neural control of rhythmic movements in vertebrates*, pages 1–40. Jon Wiley & Sons, 1988.

[Grillner *et al.* 91] S. Grillner, P. Wallén, and L. Brodin. Neuronal network generating locomotor behavior in lamprey: Circuitry, transmitters, membrane properties, and simulation. *Annu. Rev. Neurosci.*, 14:169–199, 1991.

[Grillner *et al.* 95] S. Grillner, T. Degliana, Ö. Ekeberg, A. El Marina, A. Lansner, G.N. Orlovsky, and P. Wallén. Neural networks that co-ordinate locomotion and body orientation in lamprey. *Trends in Neuroscience*, 18(6):270–279, 1995.

[Gruau 95] F. Gruau. Automatic definition of modular neural networks. *Adaptive Behavior*, 3(2):151–184, 1995.

[Ijspeert & Kodjabachian 98] A.J. Ijspeert and J. Kodjabachian. Evolution and development of a central pattern generator for the swimming of a lamprey. *In preparation*, 1998.

[Ijspeert *et al.* 97] A.J. Ijspeert, J. Hallam, and D. Willshaw. Artificial lampreys: Comparing naturally and artificially evolved swimming controllers. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life, ECAL97*, pages 256–265. MIT Press, 1997.

[Ijspeert *et al.* 98a] A.J. Ijspeert, J. Hallam, and D. Willshaw. Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology. *Dept. of Artificial Intelligence Research Paper 876, University of Edinburgh*, January 1998.

[Ijspeert *et al.* 98b] A.J. Ijspeert, J. Hallam, and D. Willshaw. Generation and investigation of the swimming circuitry of the lamprey using genetic algorithms. *In preparation*, 1998.

[Kodjabachian & Meyer 95] J. Kodjabachian and J.-A. Meyer. Evolution and development of control architectures in animats. *Robotics and Autonomous Systems*, 16:161–182, 1995.

[Kodjabachian & Meyer 98] J. Kodjabachian and J.-A. Meyer. Evolution and development of neural networks controlling locomotion, gradient-following, and obstacle-avoidance in artificial insects. *to appear in IEEE Trans. on Neural Networks*, 1998.

[Kopell 95] N. Kopell. Chains of coupled oscillators. In M.A. Arbib, editor, *The handbook of brain theory and neural networks*, pages 178–183. MIT Press, 1995.

[Lewis *et al.* 93] M.A. Lewis, A.H. Fagg, and G.A. Bekey. Genetic algorithms for gait synthesis in a hexapod robot. In Y.F. Zheng, editor, *Recent trends in mobile robots*. World Scientific, 1993.

[Lockery & Sejnowski 93] S.R. Lockery and T.J. Sejnowski. The computational leech. *Trends in Neuroscience*, 16(7):283–290, 1993.

[Matsushima & Grillner 92] T. Matsushima and S. Grillner. Neural mechanisms of intersegmental coordination in lamprey: local excitability changes modify the phase coupling along the spinal cord. *J. of Neurophysiology*, 67:373–388, 1992.

[Pearlmutter 95] B.A. Pearlmutter. Gradient calculations for dynamic recurrent neural networks: a survey. *IEEE Transactions on Neural networks*, 6(5):1212–1228, 1995.

[Schaffer *et al.* 92] J. Schaffer, D. Whitley, and L.J. Eskelman. Combinations of genetic algorithms and neural networks: a survey of the state of the art. In J. Schaffer and D. Whitley, editors, *Combinations of genetic algorithms and neural networks*. IEEE Computer Society Press, 1992.

[Viana Di Prisco *et al.* 90] G. Viana Di Prisco, P. Wallén, and S. Grillner. Synaptic effects of intraspinal stretch receptor neurons mediating movement-related feedback during locomotion. *Brain Research*, 530:161–166, 1990.

[Wadden *et al.* 97] T. Wadden, J. Hellgren, A. Lansner, and S. Grillner. Intersegmental coordination in the lamprey: simulations using a network model without segmental boundaries. *Biological Cybernetics*, 76:1–9, 1997.

[Wallén *et al.* 92] P. Wallén, Ö. Ekeberg, A. Lansner, L. Brodin, H. Traven, and S. Grillner. A computer-based model for realistic simulations of neural networks ii: The segmental network generating locomotor rhythmicity in the lamprey. *J. of Neurophysiology*, 68:1939–1950, December 1992.

[West & Wilcox 97] R.M.E. West and G.L. Wilcox. Robust parameter selection for compartmental models of neurons using evolutionary algorithms. In J.M. Bower, editor, *Computational Neuroscience: Trends in Research 1997*. Plenum Publishing, New York, 1997.

[Williams & Sigvardt 95] T.L. Williams and K.A. Sigvardt. Spinal cord of lamprey: generation of locomotor patterns. In M.A. Arbib, editor, *The handbook of brain theory and neural networks*, pages 178–183. MIT Press, 1995.

[Williams 92] T.L Williams. Phase coupling by synaptic spread in chains of coupled neuronal oscillators. *Science*, 258:662–665, 1992.

[Yao 93] Xin Yao. A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 8(4):539–567, 1993.