

Investigating Morphological Symmetry and Locomotive Efficiency using Virtual Embodied Evolution

Josh C. Bongard Chandana Paul

Artificial Intelligence Laboratory

University of Zurich

CH-8057 Zurich, Switzerland

[bongard, chandana]@ifi.unizh.ch

Abstract

The recent convergence of real-time physics-based simulation tools, the growing field of embodied cognitive science, and techniques for evolving complete agents has created a new methodology, which we refer to as Virtual Embodied Evolution. This methodology can be used to explore a wide range of issues related to the interplay between morphology and control in adaptive behaviour research. Here, we explore the intuitive, but previously unexplored correlation between morphological symmetry and locomotive efficiency in mobile, simulated agents. By evolving the morphologies and control structures of simulated agents using a genetic algorithm, it was found that agents with a higher degree of bilateral symmetry tended to exhibit greater locomotive efficiency than agents with less bilateral symmetry. This finding lends credence to the argument that for biological organisms, natural selection may have preceded, and continues to supplement sexual selection pressure favouring morphological symmetry. We conclude by discussing the future possibilities of virtual embodied evolution.

1. Introduction

The field of embodied cognitive science has developed into a coherent conceptual framework for the advancement of embodied artificial intelligence (Thelen and Smith, 1994, Clark, 1998, Pfeifer and Scheier, 1999). However, embodiment raises new research issues. Genetic and/or learning methods are often used for automating the generation of adaptive agents, and it is difficult and time-consuming to iteratively modify the shape of, and sensor and effector placements on real-world robots (Mataric and Cliff, 1996). On the other hand, developing adaptive agents completely in simulation raises its own challenges, such as effectively preserving observed behaviour of simulated agents when transferred to real-world robots

(Jakobi et al., 1995, Eggenberger et al., 1999).

One possibility for bridging the gap between simulation and the real world is by employing a physics-based simulation tool for investigating embodiment-related issues (Sims, 1994, Terzopoulos et al., 1996, Mataric et al., 1999). In this paper, the MathEngine physics-based simulation package¹ is used to study the relationship between symmetric morphology and efficient locomotion in evolved agents.

The first attempt to evolve both the morphology and control structure of simulated agents is reported in (Sims, 1994): agents were evolved for a variety of tasks using a recursive, graph-based genetic algorithm. In (Terzopoulos, 1996), a learning algorithm is used to generate behaviours for fish with three-dimensional body plans, which can deform and locomote within a simulated, physics-based environment. Ventrella (1994) also evolved morphologies for simulated agents using a genetic algorithm: initial attempts to generate symmetric morphologies by using a fitness function based solely on locomotion were not successful. Subsequent experiments built symmetry into the genotype to phenotype mapping, so that evolved agents exhibited slight variations on an underlying bilaterally symmetric body plan. However, this work did not investigate the locomotive efficiency of the evolved agents.

In biological studies, the positive correlation between morphological symmetry and locomotive efficiency has been demonstrated indirectly: it has been shown that fluctuating asymmetry (slight, random deviations from bilateral symmetry) can have an aerodynamic cost in bird species (Balmford et al., 1993, Thomas, 1993, Evans et al., 1994). In a study of the harpacticoid copepod *T. californicus*, which exhibits bilateral variation, it was found that genetic factors which influence relative limb size, in turn affecting locomotion, are expressed on both sides of the animal equally (Palmer et al., 1993).

In the biological literature, it is interesting to note that of all the reports of large-scale morphological asymmetry (Norberg, 1977, Freeman and Lundelius, 1982,

¹MathEngine PLC, Oxford, UK, www.mathengine.com

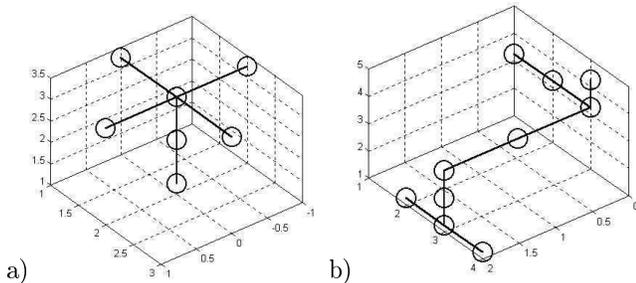


Figure 1: **Morphologies of two evolved agents.** a) shows the morphology of a symmetric agent schematically; b) shows the morphology of an asymmetric agent.

(Govind, 1989, Bock and Marsh, 1991), none of the asymmetric features investigated directly affect locomotion.

In this paper, we report a positive correlation between bilateral symmetry and locomotive efficiency for agents evolved in a physics-based, virtual task environment. Agents are evolved using two different fitness functions: one that awards for directed locomotion and bilateral symmetry, and another that awards for directed locomotion and bilateral asymmetry. We compare the locomotive efficiencies of the two types of agents.

In the next section, we describe this task environment, details of the fitness function, the genetic encoding and parameters of the genetic algorithm. In Sect. 3 we discuss the quantitative measures used for detecting locomotive efficiency. In Sect. 4 we report our results; in Sect. 5 we discuss the implications of our findings. We conclude in Sect. 6 with a discussion of the rich potential of this methodology for future studies into the interdependence of morphology and control in both simulation and for real-world embedded systems.

2. The Model

All of the agents reported here operate within a virtual, real-time physics-based environment that simulates the dynamics of multiple bodies which are affected by gravity, inertia, torque, and other internal and external forces. The morphologies of the evolved agents are treated as directed trees, similar to the agents reported in (Ventrella, 1994) and (Komosinski and Ulatowski, 1999). Each agent is composed of a number of spherical units with identical size and mass. The units are connected to each other with links of uniform length and no mass. Units can be connected to a maximum of six other units. Connections between units are constrained to the six cardinal directions up, down, north, south, east and west. Fig. 1 shows the morphologies of two agents evolved for bilateral symmetry and bilateral asymmetry, respectively.

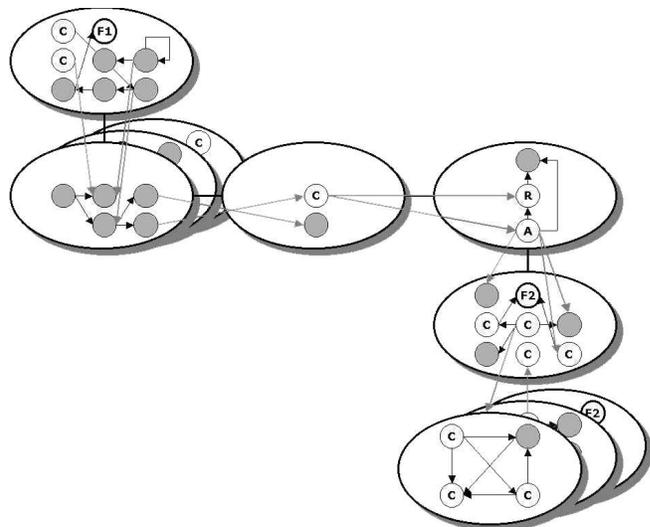


Figure 2: **A typical embedded, evolved neural network** This network was evolved to control the agent shown in Fig. 1 b): only part of the network is displayed. The darker circles F1 and F2 indicate the two types of motor neurons. The lighter circles R, A and C represent range, joint angle and contact sensors. The grey circles represent internal neurons. The large circles represent morphological units. The dark lines represent intra-unit synapses. The grey lines represent inter-unit synapses. The weights of the synapses are not shown for clarity.

2.1 Control architecture

The control of the agents is achieved through a recurrent neural network. The network is embedded within the agent's morphology. Fig. 2 shows a typical neural network, which evolved in concert with the morphology of the agent shown in Fig. 1 b). Neural connections can be constructed between connected units; neural activation to distant units can be achieved by propagating a neural signal along the synapses of neighbouring units. The neurons within the network fall into three classes: sensor neurons; motor neurons, and internal neurons. During each time step of the simulation, each neuron sums its input, applies the sigmoid activation function $\frac{1}{1+e^{-a}} - 0.5$ (where a is the summed activation to the neuron), and places the result on its output synapse(s). These results are used when the network is updated again at the next time step.

Three types of sensor neurons are used here. Contact neurons emit a maximum positive signal when the unit in which it is contained is in contact with the ground; otherwise, they emit a maximum negative signal. Proprioceptive neurons emit a signal commensurate with the current joint angle between two links connecting the parent unit and two child units; if the joint is rigid, or if the unit housing the neuron does not have two children, the neuron emits a zero signal. Range sensors emit a value inversely proportional to the distance between the unit

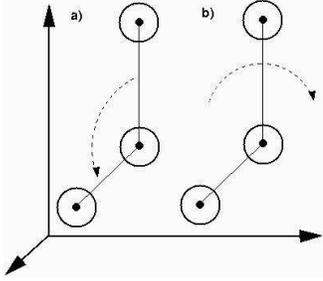


Figure 3: **The two types of joint actuation** Figs. a) and b) illustrate the different joints created by the two types of motor neurons.

housing the neuron and the single external target object in the environment. The target is placed 10 units² in the direction in which the agent should move. Thus, if the agent moves towards the target, the distance between the agent and the target will decrease, and the range sensors will emit a higher signal. By placing range sensors in different units, an agent can use a combination of differing range values to orient towards the target.

Two types of motor neurons are available for use by the agent. The presence of a motor neuron within a unit converts that unit into the central point of a one-dimensional hinge joint. The two motor neuron types correspond to the two kinds of hinge joints, with different axes of rotation (see Fig. 3).

Since the morphology is treated as a tree structure, only the first unit cannot contain motor neurons. Each unit can contain at most one motor neuron. The hinge joints are actuated using virtual springs; the elasticity and damping constants are fixed for all the agents and their constituent joints. Outputs of the motor neurons dictate changes in the equilibrium position of the virtual springs. For example, a constant, non-zero motor neuron output exerts a sequence of non-zero torques on the joint. This leads to smooth motion of the joints, regardless of whether the motor neurons emit a smooth or discontinuous signal (Pratt and Williamson, 1995).

The internal neurons can be employed by the genetic algorithm to create mappings and propagate signals between the sensor neuron inputs and the motor neuron outputs.

2.2 Genetic encoding

A variable-length genetic algorithm (Harvey, 1992) was used for evolving the agents. By using variable-length genomes, it is possible for selection pressure to evolve agents with increasing or decreasing morphological size and control structure by increasing or decreasing genome length. Initial populations of the GA contain strings of 800 bits. Selection pressure can increase this length up

²A unit in our simulation is equal to the uniform distance between any two morphological units; all other distance measures in the simulation are relative to this unit.

to a maximum of 2400 bits. Tournament selection is used, with a tournament size of three. Mutation rate is proportional to the bit string length, and performs, on average, one bit flip for each new genome generated in the population. Elitism is employed by carrying the top 50 per cent of the population into the next generation. In contrast to a developmental encoding scheme, we use a completely explicit encoding, in which each unit, connection, neuron, synapse and synapse weight directly maps onto a unique set of bits. By using a recursive rule set to grow structure, symmetric forms are more prevalent than asymmetric forms. This can be observed in the agents reported in (Sims, 1994) and (Ventrella, 1994), the symmetric neural networks grown using cellular encoding (Gruau, 1992), and the symmetric structures generated by L-systems (Rozenberg and Salomaa, 1992). Another type of developmental process, which does not contain recursive rule sets, also tends to produce symmetric structures, due to the uniform spatial distribution of transcription factors (Eggenberger, 1997).

The genome is treated as a string representation of an n-ary tree; this tree becomes the morphological frame of the agent as the read head traverses the genome. Each subset of the bit string then codes for a unit in the agent's morphology. Within this subset is contained the information necessary for constructing the local network architecture within that unit, such as the number and type of the neurons, their interconnecting synapses, and the weights of the synapses. This region also includes information for creating outgoing synapses that connect to neurons in neighbouring units. Each of the above parameters is encoded in the genome by a four-bit binary value. Fig. 4 demonstrates this mapping in more detail.

The agent's phenotype is constructed from its genotype as a read head moves linearly along the genome. Mutation or crossover sometimes adds additional bits to the end of the original genome which are insufficient for creating a new morphological unit. In such cases, the non-expressed bits are retained, in case subsequent modification reactivates this part of the genome. If genome truncation occurs instead, when the read head reaches the end of the genome, default values are supplied for the missing parameters. For example, if the last few four-bit sections of the genome are truncated, then the weights of the last few synapses in the most recently created morphological unit are not available. In this case, the weights of these synapses are set to the default value for that parameter, which is zero.

2.3 The fitness functions

Two fitness functions are used in this report: the first awards for directed movement and bilateral symmetry; the other awards for directed movement and bilateral asymmetry. The agent operates in the task environment for a specified number of simulation time steps; at the

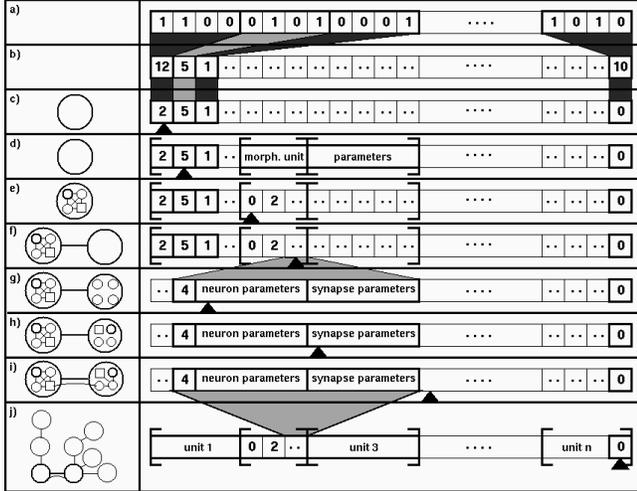


Figure 4: **The genotype to phenotype mapping.** The lefthand column shows the growth of the agent's phenotype derived from the parsing of the genotype shown in the right-hand column. Figs a) to c) show the mapping from the original bit string to a decimal, base ten representation. Fig. d) shows the placement of genetic markers for the current unit's neighbours: the first number after the start-of-unit marker indicates how many units will connect to the current unit. Fig. e) shows the creation of internal neural structure for a unit. Fig. f) shows the attachment of a neighbouring unit to a parent unit. Figs. g), h) and i) show the detailed construction of neural structure. Fig. j) shows the final phenotype of the agent reached at the end of parsing.

end of the simulation, the northern distance from the origin of the agent's southernmost unit is returned as the agent's directed movement away from the origin³.

2.4 Measuring bilateral symmetry

The bilateral symmetry of an agent is determined using the following algorithm: the vertical plane which intersects the unit whose horizontal position is closest to the average horizontal positions of all the units, is considered the plane of symmetry. The symmetry measure is then given by

$$s = \frac{4pl}{(2n - 1) - p - l}$$

where n is the total number of units comprising the agent; $2n - 1$ is the total number of units and links comprising the agent; p is the number of pairs of units lying outside the plane of symmetry, and are symmetric about that plane; and l is the number of pairs of links not contained in the plane of symmetry, and are sym-

³The southernmost unit of the agent is found by searching for the unit with a position vector containing the minimum z-component; the value of this z-component then indicates how far north the agent was able to move its trailing body part. This method for awarding directed movement eliminates the evolution of linear, passive agents, as was found in (Sims, 1994).

metric about that plane. It follows from this that agents composed of pairs of units and links which are all symmetric about the plane of symmetry attain a symmetry value of one; agents with decreasing pairs of symmetric units and links attain decreasing symmetry values; the minimum possible value is zero. Agents composed of morphological units which all fall within the plane of symmetry are given a symmetry value of zero, to avoid the evolution of two-dimensional agents: it was found that such agents produce unrealistic movement, such as tumbling motions completely within the vertical plane centred at the origin.

Thus, the two fitness functions used to evolve the agents reported here are given by ds , and $d(1 - s)$: the first awards for directed movement and bilateral symmetry; the second awards for directed movement and bilateral asymmetry.

3. Efficiency of Transport Measures

In order to compare the efficiency of transport between the symmetric and asymmetric populations, efficiency measures are used which compare the populations along axes representing different aspects of efficient locomotion.

The abstract idea of locomotive economy can be conceptualized with respect to several different criteria. In biology there is a standard nomenclature for categorizing ideas related to economy based on what variables are used (Blake, 1991). *Efficiency* is defined as performance with respect to an ideal, independent of the purpose of a task. For example, in the context of mechanics it is defined as the ratio of work or energy input to output. The other is the term *effectiveness* or *competency* in performance (Full, 1991). These definitions focus on the physical nature of a process. Effectiveness is defined as a qualitative evaluation of how a mechanism is adapted to its function. It is a study of form, and physical traits. *Perfection* is defined as the 100% efficient performance. *Optimality* represents the best performance that can be achieved given a set of limiting circumstances.

For the comparison of efficiency of transport we use three different measures related to these ideas, which together give us a robust basis for drawing qualitative conclusions about locomotive differences.

3.1 Path Efficiency

In general terms, efficiency characterizes the performance of a system relative to an ideal, applied to a single process at a time. In our simulation, every agent takes a certain path between point A , its starting point, and point B , its location at the end of the simulation. The most efficient way for the agent to travel this path is to follow the straight line between A and B . A more convoluted path between these two points indicates that the

agent’s locomotion is less efficient. The *path efficiency*, as we define it, quantitatively represents this efficiency measure. It is the ratio of the minimum distance between points A and B with respect to the length of the agent’s actual path between these points:

$$P.E. = \frac{D_{min(A-B)}}{D_{real(A-B)}}, \quad \text{where} \quad (1)$$

$$D_{min(A-B)} = \|\vec{AB}\|. \quad (2)$$

If the agent’s actual path lies exactly on the straight line from the starting point to end point, P.E. is 1, which indicates that it is 100% efficient. The further its actual path diverges from this straight line, P.E. decreases and approaches 0.

In our simulations, each agent acts for a finite time period, which is constant across simulations. However, some agents have a stochastic path with no finite periodicity, so the calculation of the absolute path efficiency is only attainable as the simulation time approaches infinity.

$$P.E.* = \lim_{x \rightarrow \infty} \frac{D_{min(A-B)}}{D_{real(A-B)}} \quad (3)$$

Our P.E. measure calculated in equation 1 is thus an approximation to this absolute efficiency and we assume that our simulation time is large enough that P.E. is asymptotically approaching the value P.E.*. It has been empirically observed that our simulation time is large enough to see large stable differences between agents’ locomotor trajectories, which supports our assumption.

3.2 Locomotive Effectiveness

Effectiveness is defined as a qualitative evaluation of how a mechanism is adapted to its purpose or function. In our simulation the agents are evolved to make the greatest possible progress in the heading direction of the target, arbitrarily defined as North. Given that the most effective way to move towards the target is to travel exactly on the straight line between the starting point and the target location T , the *Locomotive Effectiveness* quantifies the relationship between the agent’s actual path and distance moved in the target direction, as the ratio between these values.

$$L.E. = \frac{D_{North}}{D_{real(A-B)}} \quad (4)$$

$$D_{North} = \vec{AB} \cdot \vec{AT} \quad (5)$$

If the agent’s actual path lies exactly on the straight line from the starting point to the target, i.e. along vector \vec{AT} , then

$$D_{real(A-B)} = D_{North} \quad (6)$$

and the L.E. is 1, which indicates maximum effectiveness. The more its actual path diverges from this straight line, the more its L.E. value drops off.

3.3 Metabolic Efficiency

In robotics the integral over all the actuator forces is representative of the internal metabolic energy input into the system. In our model, each of the joints is actuated by a virtual damped torsional spring with spring equation:

$$F = k\theta - d\dot{\theta} \quad (7)$$

where k is the spring constant, d the damping constant, and θ the angular displacement of the spring from its equilibrium position.

The position of a joint is controlled by the motor neurons which change the equilibrium position of the joint. Thus at each time step the force applied on the arm is a function of the angular displacement between the natural angle value θ_{nat} of the joints and its actual angle, θ_{act} .

The Total Metabolic Energy (T.M.E) is a measure of the internal metabolic energy used by the agent to produce its entire sequence of motions. This can be calculated here as the integral over all the forces used by each joint:

$$T.M.E. = k \left(\int_0^\tau \theta_{act} - \theta_{nat} \right) - d(\theta_\tau - \theta_0) \quad (8)$$

Since we will only be using the T.M.E as a relative measure we choose $k = 1$ for convenience sake.

In robotics, optimality of locomotion can be measured as simply the T.M.E as defined above or as the T.M.E. value divided by the cycle period (given a periodic gait), or step length (for legged locomotion). Since our gaits may be aperiodic and without clearly identifiable steps, we use the measure of the T.M.E. value divided by the distance travelled in the target direction, D_{North} .

$$M.E. = \frac{T.M.E.}{D_{North}} \quad (9)$$

This gives us an efficiency measure in terms of energy used per unit distance and enables us to concretely compare the energy usage of agents with equal fitness.

4. Results

A total of 10 runs were performed, for 300 generations each, and using a population size of 300. After an agent is constructed from a bit string, it was allowed to act within the physics-based environment for 20,000 time steps. Five of the runs used the fitness function ds , and the other five used $d(1-s)$, where d and s are described in section 2. At the end of each run, the five most fit, unique agents were extracted from each run, and aspects of their locomotive efficiency were measured.

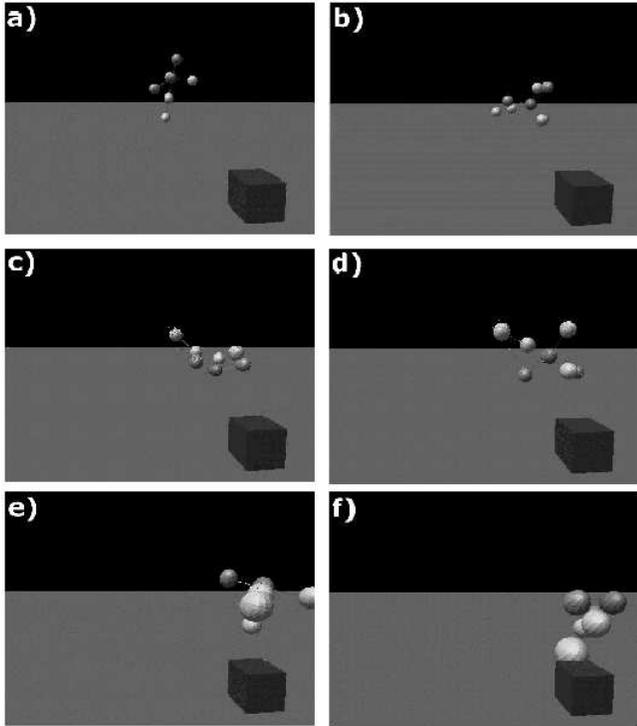


Figure 5: The motion of a symmetric agent

It was found that for the first fitness function, which awarded for movement and symmetry, the genetic algorithm rapidly converges to almost completely bilaterally symmetric (s approaches 1.0) agents. In a similar fashion, the genetic algorithm employing the fitness function awarding for movement and asymmetry rapidly converges to almost completely asymmetric (s approaches 0.0) agents. For this reason, it was possible to classify the extracted agents into two distinct classes, a symmetric and an asymmetric class.

Fig. 5 shows the behaviour of one completely bilaterally symmetric agent that was evolved. Fig. 6 shows the behaviour of an asymmetric agent. Both agents had similar fitness values. The morphologies for these agents are shown in Figs. 1 a) and b), respectively. For each evolved agent, the trajectory of its centre of mass was recorded. The trajectories of the agents shown in Figs. 5 and 6 are plotted in Fig. 7.

For each agent in the symmetric and asymmetric classes, we measured the distance travelled in the direction of the target. These distances are plotted in Fig. 8. Apart from the few symmetric agents which travel much farther than agents from either class, there is no significant difference in distance travelled.

The path efficiency of each agent, P.E. (see Eqn. 1), was calculated. Differences between the path efficiencies of the symmetric and asymmetric agents are plotted against fitness in Fig. 9.

For each of the two classes, the agents were grouped according to fitness, and D_{real} was computed for each agent, using the starting point of the agent's centre of

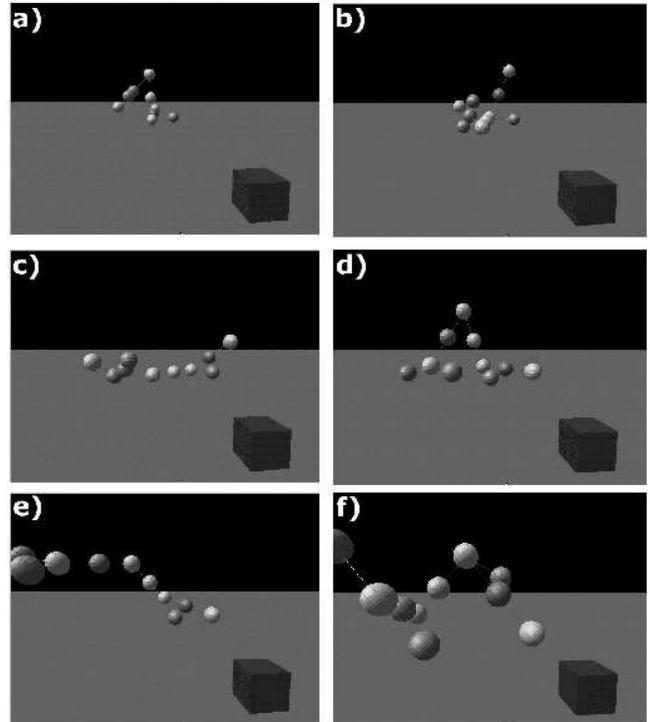


Figure 6: The motion of an asymmetric agent

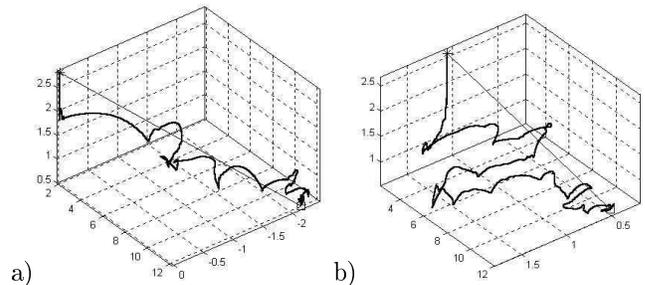


Figure 7: Trajectories for a symmetric and an asymmetric agent. Trajectories are measured as changes in the agent's centre of mass over the length of the simulation. The actual trajectories are shown using a thick line; the corresponding distance from A to B are drawn with a thin line. Note that both agents move a similar distance north, implying similar fitness values.

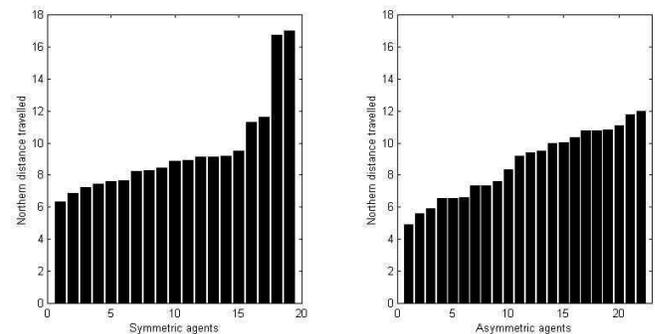


Figure 8: Distances travelled by symmetric and asymmetric agents

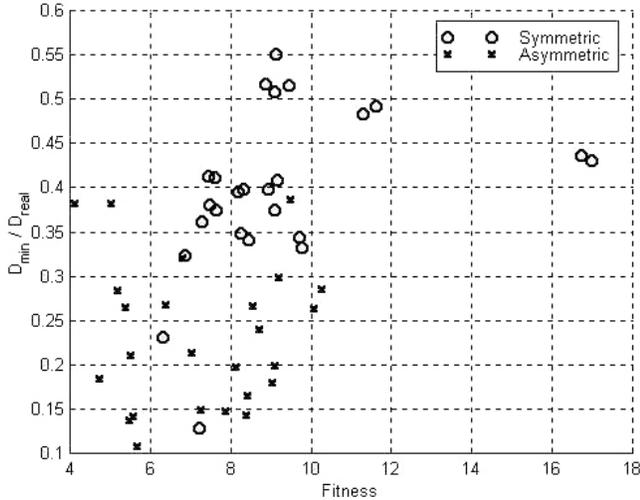


Figure 9: Path efficiencies for symmetric and asymmetric agents.

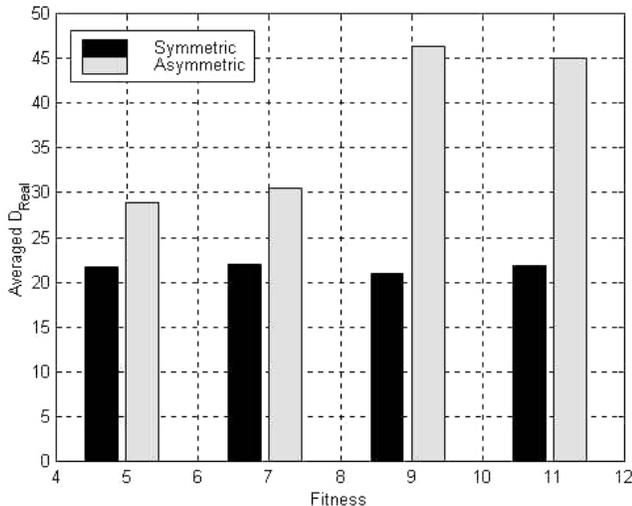


Figure 10: Differences in average, actual distance (D_{real}) travelled by agents with similar D_{north} values, indicating differences in $L.E.$

mass as point A , and the final point of its centre of mass as point B . The actual distance the agent travels between A and B is then calculated by summing the distance travelled by its centre of mass during each time step of the simulation. The D_{real} values were then averaged within each group of similarly fit agents, for both the symmetric and asymmetric agent classes. The resulting averages are shown in Fig. 10. Since $L.E.$ is defined as $\frac{D_{north}}{D_{real}}$, Fig. 10 reports differences in $L.E.$ between symmetric and asymmetric agents. In each fitness group, a lower average D_{real} value indicates that that class has a higher $L.E.$ than the other class.

The metabolic efficiency of each agent was calculated, using Eqn. 9. Agents were then grouped according to symmetry, and similar values for $M.E.$ The numbers of agents falling within these groups are shown in Fig. 11.

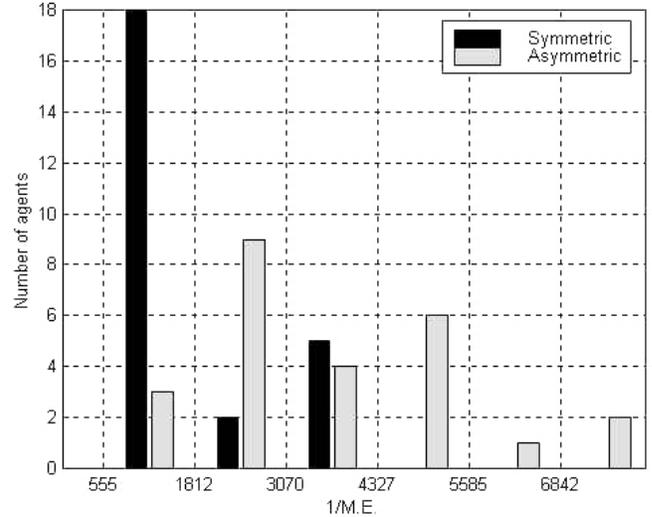


Figure 11: Differences in metabolic efficiency between symmetric and asymmetric agents. Note that the x -axis uses $\frac{1}{M.E.}$, so that agents near the y -axis have higher metabolic efficiency than agents grouped further from the y -axis.

5. Discussion

5.1 Symmetry and Efficiency

By observing the behaviours of many bilaterally symmetric and asymmetric agents, it becomes clear that the movement of asymmetric agents is almost always more erratic for asymmetric agents. The trajectories of two agents—one completely bilaterally symmetric, the other completely bilaterally asymmetric—are shown in Fig. 7. The morphologies of the two agents are shown in Fig. 1. The relative eccentricity of the asymmetric agent’s trajectory is evident from its greater deviation from the corresponding D_{min} vector.

A general trend towards greater path eccentricity for asymmetric as opposed to symmetric agents is shown by Fig. 10. For agents that travel a similar distance in the direction of the target, asymmetric agents tend to travel a further distance to reach the target than the corresponding symmetric agents. Fig. 9 shows a similar result, where the distance travelled in the direction of the target is replaced by the line-of-flight vector from the agent’s starting point to its ending point. Again, it was found that for agents with similar distances between their starting and ending points, asymmetric agents tend to travel further to achieve this distance than the corresponding symmetric agents.

In addition to lower locomotive effectiveness and path efficiency, asymmetric agents were found to be more metabolically inefficient than symmetric agents, as is shown in Fig. 11. For agents which move similar distances in the direction of the target, asymmetric agents tend to apply more total force to their actuators than corresponding symmetric agents.

5.2 Implications for Biology and Robotics

Bilateral symmetry in biological organisms is believed to have evolved only once, and has become a permanent feature of most higher animal species. However, why bilateral symmetry evolved initially is not well understood (Palmer, 1996). Also, although the prevalence of sexual selection for symmetry is widely documented (Brookes and Pomiankowski, 1994, Enquist and Arak, 1994), the origins of sexual selection for symmetry are not well explained. Our results suggest that natural selection for efficiency may be a common cause underlying both the evolution of bilateral symmetry and the origin of sexual selection for symmetry.

Initial, random variations in bilateral symmetry may have given slightly more symmetric males an evolutionary advantage due to increased locomotive or metabolic efficiency. Coupled with an initial, slight variation in female preference for symmetry, the offspring would be symmetric, and the female offspring would be both symmetric and have a higher mating preference for symmetry. Again, because symmetry implies efficiency, these symmetric females would have a selective advantage over less symmetric females and would mate more. This leads to sexual selection for symmetry: positive feedback over subsequent generations causes morphological symmetry and sexual preference for symmetry to saturate the population. In addition, due to the mechanics of sexual selection, both the preference for symmetry, and symmetry itself would become more exaggerated.

Apart from the biological implications, this work also contributes to design principles for building mobile robots. These findings support the intuition that in order to achieve directional fidelity a robot must have a near symmetric morphology. In addition, they also illuminate the less intuitive, latent correlation between symmetric morphologies and energy efficient locomotion. By making this correlation explicit, this work contributes to the central issues of efficiency in robotics research.

5.3 Morphology and Control Tradeoff

There are several ways in which physics can be exploited to achieve simplified control in agents.

An agent may exploit the physical characteristics of its morphology, such as damped springs, to create motions which do not need to be explicitly specified by the control architecture. For example, observations of several of the evolved agents' locomotion patterns have revealed that some agents exploit the physics for movement more than others. Many agents were observed to be statically unstable. These agents begin their movement by building on the momentum generated by falling forward. The behaviours are reminiscent of the techniques collectively referred to as passive dynamic control in robotics (McGeer, 1990).

Also, control can exploit the environment as a means of communicating between different parts of its morphology, reducing the need for internal communication in the control structure (Cruse et al., 1996). This second type of tradeoff is illustrated by agents which were observed to accelerate a passive joint in a forward direction by actuating a distant joint.

Evolution is able to achieve these exploitations by tuning the agent's morphology to the task. For a majority of the evolved agents with rich locomotive behaviours, the motor neurons were observed to only emit a constant signal over the length of the simulation. Although the motor neuron output is constant, the real-time interaction between the agent's control and physical dynamics produces complex behaviour. This is a clear example of how morphological adaptations can lead to reduced control complexity.

These examples illustrate that there is no positive correlation between path efficiency and metabolic efficiency. An agent with high path efficiency may be metabolically inefficient because it actuates all of its limbs over the length of the simulation. In contrast, another agent with low path efficiency may only actuate its limbs for a small fraction of the trajectory, leaving the rest to physics.

Because no positive correlation can be drawn between path efficiency and metabolic efficiency, it shows that there is no causal link between them and that they are independent measures. As it was shown that symmetric agents had both higher path efficiency and metabolic efficiency than asymmetric agents, it follows that the presence of symmetry is the cause of both kinds of efficiency.

5.4 Virtual Embodied Evolution

By evolving agents in a physics-based environment, it is possible to generate agents which are more situated and embodied than agents evolved in more abstract environments. Also, because of the increased fidelity of the simulation *vis a vis* the real world, it is easier to transport evolved designs to the real world while retaining the observed behaviour (Funes and Pollack, 1999). Therefore, it remains possible to generate and test a large number of differing body plans and related control structure completely in simulation. We refer to this methodology as Virtual Embodied Evolution.

Although some studies have reported the evolution of complete, functioning agents in a physics-based environment (Sims, 1994, Ventrella, 1994), these studies have served more as proof-of-concept investigations: several assumptions and 'tweaks' were built into the neural circuitry, genotype/phenotype mapping or morphological form in order to reduce the computational requirements, or to evolve more 'realistic' agents.

However, with the advent of commercially available physics-based simulation tools, and the continued ad-

vances in personal computer power and speed, it is now possible to use Virtual Embodied Evolution to further the maturation of concepts related to embodiment in adaptive behaviour research. This paper has investigated one such concept, namely the relationship between morphological symmetry and locomotive efficiency in evolved agents.

6. Future Research Directions

This work has focussed primarily on only one aspect of locomotion; namely, efficiency of motion in a forward direction. However, symmetry may lead to other types of locomotive economy, such as manoeuvrability. For example, it may be that a symmetric organism is better able to change direction. It seems intuitive that an asymmetric organism would exhibit unequal abilities to turn in different directions, leading to handicaps related to fleeing from predators or pursuing prey. In (Ijspeert and Kodjabachian, 1999), control architectures for swimming behaviour in simulated lampreys were investigated, and a fitness function was used which rewarded both speed and direction of motion. It is interesting to note in that work, in which a radially symmetric body plan was used, efficient locomotion was evolved.

In addition to questions related to the various types of locomotive economy, there are a host of other research questions that can be pursued with Virtual Embodied Evolution: some examples might include whether allowing selection pressure to evolve central pattern generators leads to more efficient locomotion in the resulting agents; how sensor placement affects directed locomotion; what task environments favour (or discourage) the evolution of centralized neural structure; or how the addition of various developmental mechanisms to an explicit genotype/phenotype mapping (such as the one presented here) affect the convergence to fit agents in the genetic algorithm.

7. Conclusion

Through the use of an explicit genotype/phenotype mapping, which does not implicitly favour either morphological or control symmetries, distinct sets of bilaterally symmetric and asymmetry agents were evolved by using two fitness functions, one which awards for locomotion and symmetry, and the other for locomotion and asymmetry. It was then shown, by comparing a suite of efficiency measures against morphological symmetry, that evolved agents with relatively high bilateral symmetry tend to move more efficiently than highly asymmetric agents.

The result that bilateral symmetry leads to locomotive and metabolic efficiency in the evolved agents reported here suggests that there may be a common cause underlying the evolution of bilateral symmetry and sexual

selection for symmetry. It is hoped that this work will lead to more biological investigations into these issues.

This work has made explicit the connections between physics-based simulation, increased computing power, the maturation of concepts in embodied cognitive science, and evolutionary techniques. This confluence of ideas is referred to as Virtual Embodied Evolution, which represents a unique methodology for studying adaptive behaviour.

8. Acknowledgements

Many thanks to Matt Hare for his insightful contributions to this work, and Torsten Reil, for his thoughtful commentary. The authors would also like to thank Gabriel Gómez, Massimiliano Lungarella and Verena Hafner for their spare CPU cycles.

References

- Balmford, A., Jones, I. L., and Thomas, A. L. R. (1993). On avian asymmetry: evidence of natural selection for symmetrical tails and wings in birds. *Procs. of the Royal Society of London B*, 252:245–251.
- Blake, R. W., (Ed.) (1991). *Efficiency and economy in animal physiology*. Cambridge University Press, Cambridge, UK.
- Bock, G. R. and Marsh, J., (Eds.) (1991). *Biological Asymmetry and Handedness*. Wiley and Sons, New York, NY.
- Brookes, M. and Pomiankowski, A. (1994). *Trends Ecol. Evol.*, 9:21–25.
- Clark, A. (1998). *Being There: Putting Brain, Body, and World Together Again*. Bradford Books, Cambridge, MA.
- Cruse, H., Bartling, C., Dean, J., Kindermann, T., Schmitz, J., Schumm, M., and Wagner, H. (1996). Coordination in a six-legged walking system: simple solutions to complex problems by exploitation of physical properties. In P. Maes, M. J. Mataric, et al., (Ed.), *Procs. of the Fourth Intl. Conf. on the Simulation of Adaptive Behavior*, pages 84–93. MIT Press.
- Eggenberger, P. (1997). Evolving morphologies of simulated 3d organisms based on differential gene expression. *Procs. of the Fourth European Conf. on Artificial Life*, pages 205–213.
- Eggenberger, P., Ishiguro, A., Tokura, S., Kondo, T., Kawashima, T., and Aoki, T. (1999). Toward seamless transfer from simulated to real worlds: a dynamically-rearranging neural network approach.

- Eighth European Workshop on Learning Robots*, pages 4–13.
- Enquist, M. and Arak, A. (1994). Symmetry, beauty and evolution. *Nature*, 372:169–172.
- Evans, M. R., Martins, T. L. F., and Haley, M. (1994). The asymmetrical cost of tail elongation in red-billed streamertails. *Procs. of the Royal Society of London B*, 256:97–103.
- Freeman, G. and Lundelius, J. W. (1982). The developmental genetics of dextrality and sinistrality in the gastropod *lymnaea peregra*. *Wilhelm Roux's Archives of Developmental Biology*, 191:69–83.
- Full, R. J. (1991). The concepts of efficiency and economy in land locomotion. In Blake, R. W., (Ed.), *Efficiency and Economy in animal physiology*, pages 97–132. Cambridge University Press.
- Funes, P. and Pollack, J. (1999). Computer evolution of buildable objects. In Bentley, P., (Ed.), *Evolutionary Design by Computer*, pages 387–403. Morgan Kaufman, San Francisco.
- Govind, C. K. (1989). Asymmetry in lobster claws. *American Scientist*, 77:468–474.
- Gruau, F. (1992). Cellular encoding of genetic neural networks. Technical Report 92-21, Ecole Normale Supérieure de Lyon, Institut IMAG.
- Harvey, I. (1992). Species adaptation genetic algorithms: A basis for a continuing saga. *Procs. of the First European Conf. on Artificial Life*, pages 346–354.
- Ijspeert, A. J. and Kodjabachian, J. (1999). Evolution and development of a central pattern generator for the swimming of a lamprey. *Artificial Life*, 5(3):247–269.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: the use of simulation in evolutionary robotics. *Procs. of the Third Intl. Conf. on Artificial Life*, pages 704–720.
- Komosinski, M. and Ulatowski, S. (1999). Framsticks: Towards a simulation of a nature-like world, creatures and evolution. *Procs. of the Fifth European Conf. on Artificial Life*, pages 261–265.
- Mataric, M. J. and Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1):67–83.
- Mataric, M. J., Zordon, V. B., and Williamson, M. M. (1999). Making complex articulated agents dance: an analysis of control methods drawn from robotics, animation and biology. *Autonomous Agents and Multi-Agent Systems*, 2(1):23–44.
- McGeer, T. (1990). Passive dynamic walking. *Intl. Jour. Rob. Res.*, 9:62–82.
- Norberg, R. A. (1977). Occurrence and independent evolution of bilateral ear asymmetry in owls and implications on owl taxonomy. *Phil. Trans. of the Royal Soc. of London B*, 280:375–408.
- Palmer, A. R. (1996). From symmetry to asymmetry: Phylogenetic patterns of asymmetry variation in animals and their evolutionary significance. *Proc. Natl. Acad. Sci. USA*, 93:14279–14286.
- Palmer, A. R., Strobeck, C., and Chippindale, A. K. (1993). Bilateral variation and the evolutionary origin of macroscopic asymmetries. *Genetica*, 89:201–218.
- Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. MIT Press, Cambridge, MA.
- Pratt, G. A. and Williamson, M. M. (1995). Series elastic actuators. *IEEE Intl. Conf. on Intelligent Robots and Systems*, 1:399–406.
- Rozenberg, G. and Salomaa, A., (Eds.) (1992). *Lindenmeyer Systems : Impacts on Theoretical Computer Science, Computer Graphics, and Developmental Biology*. Springer-Verlag, Berlin.
- Sims, K. (1994). Evolving 3d morphology and behaviour by competition. *Artificial Life IV: Procs. of the Fourth Intl. Workshop on the Synthesis and Simulation of Living Systems*, pages 28–39.
- Terzopoulos, D., Rabie, T., and Grzeszczuk, R. (1996). Perception and learning in artificial animals. *Artificial Life V: Procs. of the Fifth Intl. Conf. on the Synthesis and Simulation of Living Systems*, pages 313–320.
- Thelen, E. and Smith, L. B., (Eds.) (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press, Cambridge, MA.
- Thomas, A. L. R. (1993). The aerodynamic cost of asymmetry in the wings and tail of birds: asymmetric birds can't fly round tight corners. *Procs. of the Royal Society of London B*, 254:181–189.
- Ventrella, J. (1994). Explorations of morphology and locomotion behaviour in animated characters. *Artificial Life IV*, pages 436–441.