The Road Less Travelled: Morphology in the Optimization of Biped Robot Locomotion

Chandana Paul and **Josh C. Bongard**

Artificial Intelligence Laboratory
University of Zurich
CH-8057 Zurich, Switzerland
[chandana,bongard]@ifi.unizh.ch

Abstract

In this paper, stable bipedal locomotion has been achieved using coupled evolution of morphology and control on a 5-link biped robot in a physics-based simulation environment. The robot was controlled by a closed loop recurrent neural network controller. The goal was to study the effect of macroscopic, midrange and microscopic changes in mass distribution along the biped skeleton to ascertain whether optimal morphology and control pairs could be discovered. The sensor-motor coupling determined that small changes in morphology manifest themselves as large changes in the performance of the biped, which were exploited by the optimization process. In this way, mechanical design and controller optimization were reduced to a single process, and more mutually optimized designs resulted. This work points to alternative routes for efficient automated and manual biped optimization.

1 Introduction

In the traditional approach to biped design and optimization, designers usually preselect a morphology. They then design a controller for the preselected mechanical design. Often, the biped and the controller are implemented in simulation to predetermine if together they lead to the desired performance, before the actual robot is built. Finally, the robot is built in the real world, and the controller parameters are tuned. This methodology has led to many successful biped walking robots such as the Spring Flamingo [Pratt and Pratt, 1998], the Honda Humanoid [Hirai *et al.*, 1998] and the BIP2000 [Espiau and Sardain, 2000]. The success of this methodology has encouraged robot designers to design morphologies first and then pursue further performance improvements with controller enhancements. This methodology has also been adopted in evolutionary robotics. There are numerous examples of biped controller design and optimization using genetic algorithms, where the biped morphology has been preselected and fixed. For example, [Fukuda *et al.*, 1997] use a GA for stabilization control of a simulated fixed morphology biped robot. Similarly, in work by [Reil, 1999] a GA is used to evolve the weights of a recurrent neural controller for walking in a fixed simulated biped.

The strong focus on control has left a second avenue, that of *morphology* enhancements leading to large performance increases, relatively unexplored. In recent times, advances in Embodied AI have increasingly stressed the importance of the robot body [Brooks and Stein, 1994] and morphological parameters, such as mechanical structure and sensor and actuator placements, which greatly influence the performance of a robot [Pfeifer and Scheier, 1999]. More significantly, work in passive dynamic walking by [McGeer, 1990] has shown that morphology is all important in determining the performance of a passive bipedal mechanism. However, this effect of morphology has not been systematically studied or harnessed to tune performance in the domain of active biped robot locomotion.

Explorative studies of the coevolution of morphology and control have addressed the issue indirectly in the past. Pioneering work on the evolution of morphologies and neural controllers was conducted in a virtual physics based simulation environment by [Sims, 1994] and made clear the strong interdependence between morphology and control. Similar work in coevolving controllers and body plans for obstacle avoidance in simulation followed by [Lee *et al.*, 1996] and for real world robots by [Lipson and Pollack, 2000]. The effect of morphological attributes specifically on locomotive performance was studied by [Bongard and Paul, 2000] and [Chocron and Bidaud, 1999]. While these projects provide the theoretical inspiration for this work, their loosely constrained morphologies prevent any direct association with issues in biped locomotion.

More recently, this methodology of jointly optimizing morphology and control was applied to develop a pseudo biped robot, in which the two moving legs driven by open loop oscillators were stabilized by a long flat tail [Juárez-Guerrero *et al.*, 1997]. However, the effects of morphology on control have not yet been truly integrated into the primary design process of more anthropomorphic biped robots. In this paper, for the first time morphological parameters and a closed loop controller have been optimized simultaneously for a biped robot, so that the mechanical and controller design are carried out in a single process. To demonstrate the use of this methodology, we address the engineering design problem of mass distribution along the biped skeleton, one of the first

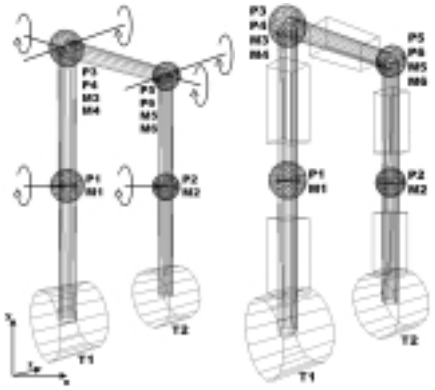

Figure 1: Biped construction: Fig. a) shows the biped skeleton and its 6 degrees of fredom Fig. b) shows the biped with attached mass blocks

mechanical design decisions. The distribution of the mass is represented as a problem of parametrically positioning discrete blocks along the length of the biped, and determining their geometrical dimensions. This problem was chosen because it was analogous to designer decisions regarding relative positions of motors and gears, usually the heaviest components, along the biped skeleton. Thus the goal was to combine the basic biped design and controller optimization process, so that the controller could take advantage of the effect if changing mass distributions on the dynamics of the biped. We show that this is not only a successful method for finding stable biped locomotion but also for taking advantage of the intrinsic morphology and control tradeoff, which ensures the development of stable gaits.

2 The Robot

The robot is a 5-link biped robot with 6 degrees of freedom, simulated in a real-time, physics-based virtual environment[1]. The robot has a waist, two upper legs and two lower leg links as shown in Fig 1. Each knee joint, connecting the upper and lower leg links, has one degree of freedom in the sagittal plane. Each hip joint, connecting the upper leg to the waist, has two degrees of freedom: one in the sagittal plane and one in the frontal plane. These correspond to the roll and pitch motions.

The joints are limited in their motion with joint stops, with ranges of motion closely resembling those of human walking. The hip roll joint on each side has a range of motion between $-\pi/7$ and $\pi/7$ degrees with respect to the frontal plane. The hip pitch joint has a range of motion between $-\pi/10$ and $\pi/10$, with respect to the sagittal plane. The knee joint has a range of motion between $-\pi/2$ and 0 with respect to the axis of the upper leg link to which it is attached.

Each of the joints is moved by a simulated torsional actuator. The actuator receives position commands from the controller. It uses proportional control to determine the velocity of the link, with a relatively low maximum torque ceiling.

[1]MathEngine PLC, Oxford, UK, `www.mathengine.com`

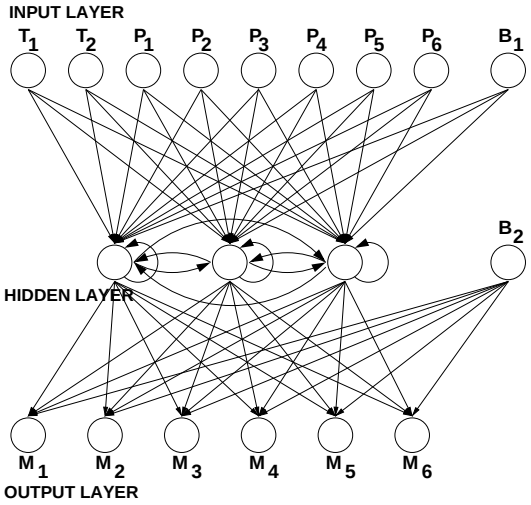

Figure 2: gives a pictorial representation of the neural network used to control both types of agents. `T1` and `T2` correspond to the two touch sensors, `P1` through `P6` indicate the six proprioceptive sensors, and `M1` through `M6` indicate the six torsional motors of the biped. `B1` and `B2` indicate the two bias neurons included in the network.

The torque applied to actuators is determined by

$$\tau_{t+1} = max(I(\omega_t - k(\theta - \theta_d)), \tau_{max}) \quad (1)$$

where θ is the actual joint angle, θ_d is the desired joint angle, τ_{max} is maximum torque ceiling, $\omega = \dot{\theta}$, and I is the inertia matrix.

This means that the velocity of a link will be greater the further it is from the commanded joint angle position, but if the force required to achieve this velocity is too large, it will only apply the maximum force. This mechanism incorporates a measure of compliance into the system, and is in accordance with the capabilities of real world actuators.

3 Neural Controller

The neural controller used is a recurrent neural network, which provides the intrinsic capability of producing cyclic dynamics. The choice of this network is based on the work of [Reil, 1999] and [Gallagher *et al.*, 1996]. The agent contains two haptic sensors in the feet, and six proprioceptive sensors and torsional actuators attached to the six joints, as outlined in Figs. 1 a) and b). At each time step of the simulation, agent action is generated by the propagation of sensory input through the network shown in Fig. 2 and the values of the output layer are fed into the actuators as desired positions.

The input layer contains nine neurons, with eight corresponding to the sensors, and an additional bias neuron. All neurons in the network emit a signal between -1 and 1: the haptic sensors output 1 if the foot is in contact with the ground, and -1 otherwise; the proprioceptive sensor values are scaled to the range $[-1, 1]$ depending on their corresponding joint's range of motion; and bias neurons emit a constant signal of 1. The input layer is fully connected to a hidden layer composed of three neurons. The hidden layer is fully

Table 1: Morphological parameters of the agent. A `ul` represents one unit length defined as the radius of the spherical sockets at the hip and knees. A `um` represents one unit mass defined as the mass of the same spherical socket. The parameters in boldface vary under evolutionary control between the ranges shown below.

Index	Object	Dimensions	Mass
1	Knees	r = 1 ul	1 um
2	Hip sockets	r = 1 ul	1 um
3	Feet	r = 2 ul, w = 3 ul	1 um
4	Lower Legs	r = 0.5 ul, h = 8 ul	0.25 um
5	Upper Legs	r = 0.5 ul, h = 8 ul	0.25 um
6	Waist	r = 0.5 ul, w = 8 ul	0.25 um
7	Waist Mass Block	**l = [0.4, 3.6] ul, w = h = [0.2,3.0] ul**	**L = [0.001, 0.246]um M = [0.002, 0.505]um H = [0.004, 1.009]um**
8	Lower Mass Block	**l = [0.4, 3.6] ul, w = h = [0.2,3.0] ul**	**L = [0.001, 0.246]um M = [0.002, 0.505]um H = [0.004, 1.009]um**
9	Upper Mass Block	**l = w = [0.2,3.0] ul, h = [0.4, 3.6] ul**	**L = [0.001, 0.246]um M = [0.002, 0.505]um H = [0.004, 1.009]um**
Index	Joint	Plane of Rotation	Range of Motion
10	Knee	sagittal	$-\frac{\pi}{2} \rightarrow 0$ (radians)
11	Hip	sagittal	$-\frac{\pi}{7} \rightarrow \frac{\pi}{7}$
12	Hip	frontal	$-\frac{\pi}{10} \rightarrow \frac{\pi}{10}$

and recurrently connected, plus an additional bias neuron. The hidden and bias neurons are fully connected to the eight neurons in the output layer. The activations of the hidden and output neurons are computed by

$$a = \sum_{i=1}^{n} w_{ij} O_i \tag{2}$$

where O_i is the output of a neuron in the previous layer, and w_{ij} is the weight of the synapse connecting them. The output of this neuron is then given by

$$O = \frac{2}{1 + e^{-a}} - 1 \tag{3}$$

The values at the output layer are scaled to fit the range of their corresponding joint's range of motion. Torsion is then applied at each joint to attain the desired joint angle.

4 The Genetic Algorithm

A fixed length genetic algorithm was used to evolve the controllers reported in this paper. Each run of the genetic algorithm was conducted for 300 generations, using a population size of 300.

At the end of each generation, the 150 most fit genomes were preserved; the others were deleted. Tournament selection with a tournament size of three, is employed to probabilistically select genotypes from among those remaining for mutation and crossover. 38 pairwise one-point crossings produce 76 new genotypes: the remaining 74 new genotypes are mutated copies of genotypes from the previous generation. The mutation rate was set to generate an average of five mutations for each new genome created. Mutation involved the replacement of a single value with a new random value.

Each genome contains floating-point values encoding the 60 synaptic weights of the neural network, plus any additional morphological parameters. These values are rounded to two decimal places and range between -1.00 and 1.00.

Each individual is evaluated for 2000 runs of the simulation. The evaluation is prematurely terminated if the center of gravity of the waist drops below the original vertical position of its knees (it falls) or if both feet lift off the ground (it starts to run). This second termination criteria was added because the primary interest of this project was to study walking, and not running gaits. At the end of the evaluation, the distance of the biped travelled in the sagittal plane (determined relative to its original position) was considered its fitness.

5 Experiments

In order to study the effect of a varying range of mass distributions on the performance of the biped three sets of experiments were conducted. In each of the three experiments, the total mass of the blocks distributed along the skeleton represented a different fraction of the total mass of the biped. We wanted to study the effect of macroscopic, midrange and microscopic changes in weight distribution of the biped skeleton to study whether optimal morphology/control pairs could be discovered over a range of mass distributions. The lengths, widths and vertical attachment points of the blocks of the lower and upper leg, and the length and width of the block on the waist were supplied as the morphological parameters

subject to evolutionary control. The block on the waist was always fixed at the midpoint of the waist to maintain bilateral symmetry. The rectangular blocks on the legs could be as long as the leg link to which they were attached, and were not allowed to extend beyond the endpoint of the link. There were always two blocks per leg, each of which could be attached to the upper or lower leg link. The size and mass ranges for the blocks are summarized in Table 1. Although the sizes of the blocks varied, the total mass of the blocks added together was kept constant. Thus the goal was to study the effect of mass distribution of the same total mass along the body.

In the first set of experiments with microscopic changes in mass distribution, the blocks were allowed the same range of geometric variability as described above but their total mass was normalized to 0.25 um. Thus, the mass of the blocks represented a 6% of the total mass of the biped. In this second set of experiments with mid-range changes, the blocks were now normalized to a total weight of 0.513 um. This represented 13% of the total weight of the biped. In the final set with macroscopic changes, the total mass of the blocks was 1.025um, representing 27.3% of the total mass of the biped. In each of the three cases, 20 evolutionary runs were conducted, one run last approximately 40 minutes. Thus the results from 40 hours of data collection are presented below.

6 Results

6.1 Microscopic Changes

In the case of microscopic changes in weight distribution, 20 runs were performed. Of these runs, only one evolved stable locomotion (See Fig 3). This is not a surprising result, as the problem of achieving a stable limit cycle in bipedal locomotion is not trivial. The best biped in this set of experiments had a morphology optimized to the configuration as shown in Fig. 4, which was largely different from the initial setting of its morphology. Both leg blocks were fixed to the upper leg link. One block was thicker and shorter and was positioned higher on the leg, and the other narrower but longer was positioned a bit lower on the same link. The waist block is also considerably large, leading to a vertical center of mass at 0.542. The gait of this biped demonstrated a rapid shuffling motion.

6.2 Mid Range Changes

In the case of mid-range mass distribution changes, 20 runs were also performed. Of these runs, 2 evolved stable locomotion (See Fig. 5). The maximum fitness attained by the best biped was 44.2. This biped had a morphology as shown in Figure 6, which was also largely different from its initial setting. In this case again, both the leg blocks are positioned on the upper leg. One block is thick and long, hiding the other block inside it which is thinner and a bit shorter. The waist mass is also considerably sizeable.

6.3 Macroscopic Changes

Finally, in the case of macroscopic changes to the mass distribution, 20 runs were also performed. Of these three evolved stable locomotion. The best fitness achieved in this case was 32.1 as shown in Figure 7. This run produced the agent shown

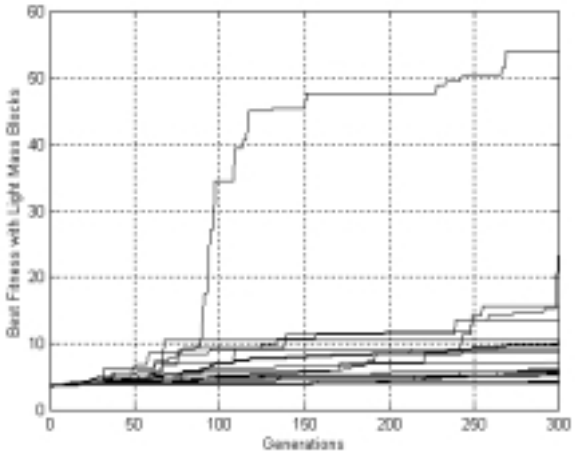

Figure 3: Best fitness achieved in each generation with microscopic mass distribution changes

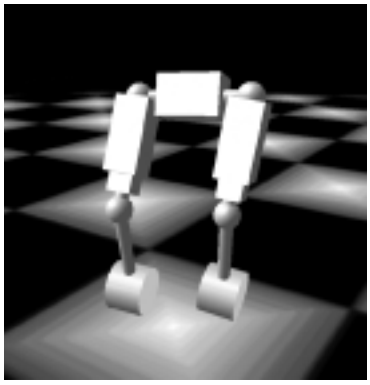

Figure 4: Morphology of most fit biped achieved with microscopic mass distribution changes

in Figure 8. In this case, the waist mass has been reduced. Conversely, the upper legs on which the two blocks are again positioned (one inside the other), are even heavier.

7 Discussion

Stable Locomotion Using Morphology
The experimental results have shown for the first time that stable locomotion can be implemented using coupled optimization of morphology and control in biped robots. Six successful examples of this process have been charted here. This shows that the success of our methodology is not infrequent or difficult to achieve. The results also show that this methodology is robust in the face of varying ranges of morphological parameter settings.

Morphology for Performance Optimization
In each of the three cases shown, the final morphology of the biped is largely different from the configuration of the best biped at the beginning of the run. However, although the range of potential variability in the size and placement of the blocks and the vertical position of the center of mass was

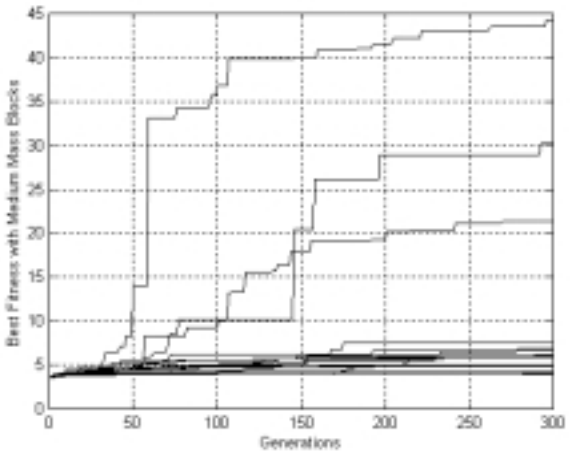

Figure 5: Best fitness achieved in each generation with mid-range mass distribution changes

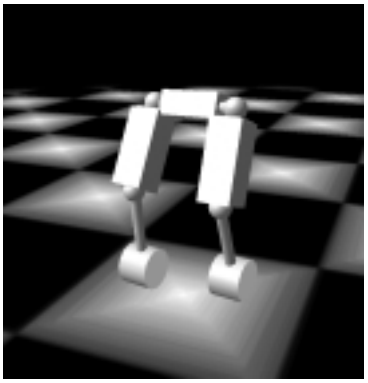

Figure 6: Morphology of most efficient biped achieved with mid-range mass distribution changes

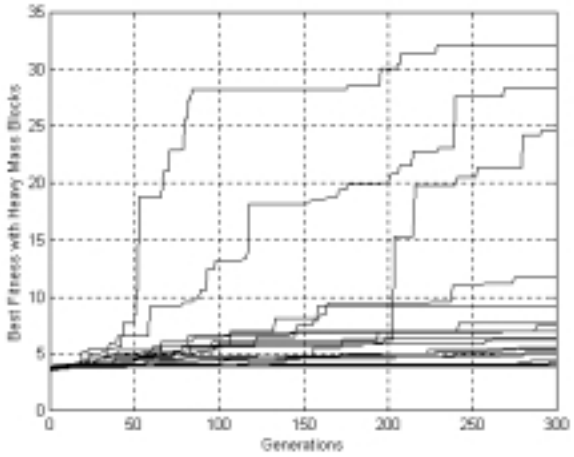

Figure 7: Best fitness achieved in each generation with macroscopic changes in mass distribution

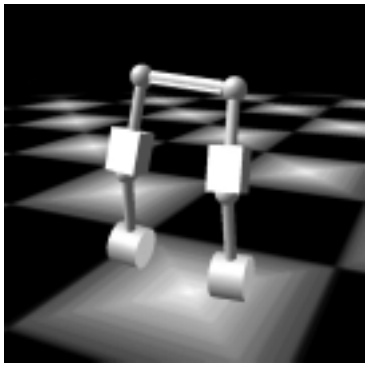

Figure 8: Morphology of the most fit biped achieved with macroscopic changes in mass distribution

large (between 0.4 and 0.6), the morphologies of the three best walkers all had no blocks on their lower legs, and their vertical CoM positions fell in a narrow band between 0.54 and 0.56 (Fig. 9). This suggested that within the space of morphological variability a certain region was more likely to yield stable locomotion than others and that the morphology was not entirely inconsequential.

Comparing the three graphs of best fitnesses in Figures 3, 5 and 7, another significant difference is observed. In the microscopic case, there is only one case of stable locomotion. In the mid-range case, two cases of stable locomotion are seen and in the macroscopic case, three cases were discovered. Unless these results are purely coincidental, it suggests that the greater the possibility for morphological change, the greater is the likelihood of achieving stable locomotion.

While these results indirectly suggest the importance of morphology in performance optimization, more concrete evidence was collected: in some of the runs, the trajectories of the centers of mass of the bipeds were tracked during their evaluation. In several of these, it was seen that morphological mutations significantly improve performance. In one population, a mid-range biped was tracked during its life history: its center of mass reached a distance of 11.1. It was replaced by its child, which sustained eight point mutations, and its center of mass travelled in 20.8 ul in the desired direction. The trajectories of the centers of mass of these two agents are indicated in Fig. 10 by the light gray and dark gray lines, respectively. Of the eight mutations, one of these was a morphological change. A third agent was tested, which was genotypically equivalent to the more fit child, except that the morphological mutation was suppressed. This third agent's center of mass reached a distance of 15.6 ul, and its trajectory is indicated by the black line in Fig. 10. The control parameter mutations serve to stabilize the gait and to straighten the biped trajectory to some extent. The morphological mutation further corrects the robot's direction of travel, and increases its total distance travelled. It is thus clear from the results that morphological change can be used by an optimization algorithm to tune robot performance.

It is no surprise that this should be the case. In a closed loop system such as the one implemented here, the controller recieves sensory input from the biped mechanical system, such as joint angle and foot contact information. This sensory input is processed through the network, and position

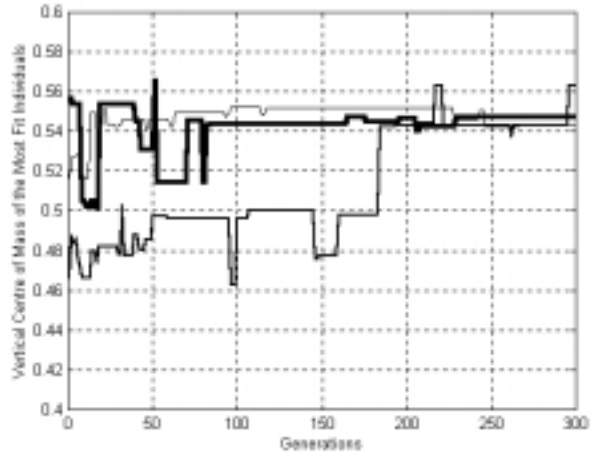

Figure 9: Changes in vertical position of biped Center of Mass: The thin line tracks the changes of the vertical CoM for the most fit biped in the microscopic case. The middle line tracks the CoM of the best biped in the mid-range case and the thick line is that of the best biped in the macroscopic case.

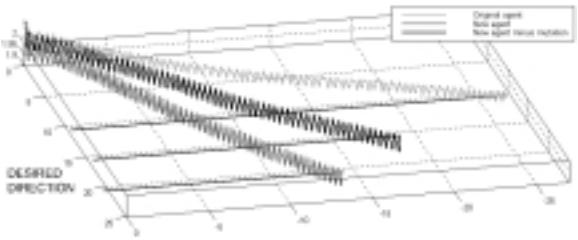

Figure 10: Trajectories of CoM of three Bipeds: The light gray trajectory shows the performance of an evolved biped. The dark gray line shows the performance of one of its children, which has seven control and one morphological mutation. The black trajectory shows the performance of a third biped, which is equivalent to the more fit child, except that the morphological mutation is suppressed.

commands are sent to the motors. The position control is not precise and rigid, but includes a large measure of compliance so that small mechanical changes affect the actual joint angle positions, with more severe changes affecting foot contact. This in turn affects the next set of sensory inputs to the controller and the next set of motor commands to the biped. Thus the controller and the morphology form an inseparable closed loop dynamical system, in which a slight change to any parameter can have large repercussions on the overall behavior of the system. What is surprising, however, is that this effect has not been utilized more effectively in biped robot optimization.

Performance vs. Efficiency
The results here show concretely that morphological parameters *can* be used to optimize performance in a biped robot. What they do not necessarily suggest is that the performance improvement gained from a morphological change will be larger than that of a controller change. The consideration of a morphological change is therefore not based on performance, but on efficiency. Sometimes, making a simple change in weight distribution on a biped robot (by attaching weights, for example) can be much more efficient than reprogramming a controller, especially when the very same controller would suffice with a slightly altered morphology. Thus, including morphology in the optimization process can lead to much more efficient design in some cases.

Combined Mechanical and Controller Design
In addition to the fact that including morphology in the optimization process leads to greater efficiency, if included from the beginning of a design project it can also help to conflate the two phases of robot design into one process. When mechanical design considerations, such as the placement of motors and gears in this case, are included in the combined optimization process, the end products are not only an optimized controller, but also a prototype mechanical design which is highly tuned to the controller.

Implications for Future Biped Design
It is hoped that this research will lead the way to further investigations into the use of morphology for optimizing robot performance. While this research uses a genetic algorithm for automated robot design, the implications of this work are by no means limited to GAs or to automated robot design at all. The broader implications of this research point to the possibilities for optimization of biped robot performance through the morphology. Currently, when biped robot performance does not meet designer expectations, engineers and programmers spend all their effort implementing controller changes until the robot's performance is as desired. It is at this juncture that we would like to encourage engineers to consider the alternative route of making morphology modifications. As has been shown in this work, often a simple morphology modification can lead to a large increase in performance, equivalent to several man hours of reprogramming. While it may be argued that such morphological changes to a robot may be much more work intensive than software alterations, this is simply because biped robots in the past have been built with preconceived morphological finality and have been very difficult to alter *post hoc*. But this need not be the case. If future biped designs allow for greater morphological flexibility, more alternative routes to successful bipeds could be discovered during the optimization process.

The insights gained from this methodology will be used for the design of a real physical biped robot at the Artificial Intelligence Lab, University of Zurich. The biped will have 8 degrees of freedom, and the simulation prototype will be extended to include additional parameters to further constrain the robot design. This will ensure that the biped is optimized to perform its basic locomotion with a simple neural network controller architecture similar to that shown here. Additional morphological parameters will be included in the optimization of robustness of the robot against both non-level walking surfaces and external disturbances.

8 Conclusions

This paper outlines a new approach to the optimization of biped robot locomotion: inclusion of morphological parameters in the optimization process. We describe three sets of experiments in which morphological parameters pertaining to mass distribution are considered coupled with control parameters in the evolution of stable bipedal gaits. The methodology shows how mechanical design decisions and controller optimization can be accomplished in a single step and can lead to more mutually optimized systems. This work opens up the possibility for robot designers to make iterative or systematic *post hoc* morphological changes to further tune biped performance.

References

[Bongard and Paul, 2000] Josh C. Bongard and Chandana Paul. Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. In *The Sixth International Conference on the Simulation of Adaptive Behaviour*, pages 420–429, Paris, France, August 2000.

[Brooks and Stein, 1994] R. A. Brooks and L. A. Stein. Building brains for bodies. *Autonomous Robots*, 1(1):7–25, 1994.

[Chocron and Bidaud, 1999] O. Chocron and P. Bidaud. Evolving walking robots for global task based design. In *Proceedings of the Congress on Evolutionary Computation*, pages 405–412, Washington, USA, June 1999.

[Espiau and Sardain, 2000] B. Espiau and P. Sardain. The anthropomorphic biped robot BIP2000. *Proceedings of the 2000 IEEE Conference on Robotics and Automation*, 4:3996–4001, 2000.

[Fukuda *et al.*, 1997] T. Fukuda, Y. Komota and T. Arakawa. Stabilization control Of biped locomotion robot based learning with GAs having self-adaptive mutation and recurrent neural networks. *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 217–222, 1997.

[Gallagher *et al.*, 1996] J. C. Gallagher, R. D. Beer, K. S. Espenschied and R. D. Quinn. Application of evolved locomotion controllers to a hexapod robot. *Robotics and Autonomous Systems*, 19:95–103, 1996.

[Hirai *et al.*, 1998] K. Hirai, M. Hirose, Y. Haikawa and T. Takenaka. The development of the Honda humanoid robot. In *Proceedings of the IEEE Conference on Robotics and Automation*, Leuven, Belgium, April 1998.

[Juárez-Guerrero *et al.*, 1997] S. Juárez-Guerrero, S. Muñoz-Gutiérrez and W. W. Mayol-Cuevas. Design of a walking machine structure using evolutionary strategies. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, pages 1427–1432, San Diego, California, 1997.

[Lee *et al.*, 1996] W. Lee, J. Hallam and H. H. Lund. A hybrid GP/GA approach for co-evolving controllers and robot bodies to achieve fitness-specific tasks. In *Proceedings of IEEE Third International Conference on Evolutionary Computation*, Nagoya, Japan, 1996.

[Lipson and Pollack, 2000] H. Lipson and J. B. Pollack. Automatic design and manufacture of artificial lifeforms. *Nature*, 406:974–978, 2000.

[McGeer, 1990] T. McGeer. Passive dynamic walking. *Int. J. Robotics Research*, 9(2):62–82, 1990.

[Pfeifer and Scheier, 1999] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, Cambridge, Massachusetts, 1999.

[Pratt and Pratt, 1998] J. Pratt and G. Pratt. Exploting natural dynamics in the control of planar bipedal walking robot. In *Proceedings of the 36th Annual Allerton Conference on Communication, Control and Computing*, Monticello, IL, 1998.

[Reil, 1999] T. Reil. Evolving bipedal locomotion in a physical simulator. Masters Thesis, University of Sussex, School of Cognitive Sciences, Brighton, UK, 1999.

[Sims, 1994] K. Sims. Evolving 3D morphology and behaviour by competition. In *Artificial Life IV*, pages 28–39, Cambridge, MA, 1994.