

Making Evolution an Offer It Can't Refuse: Morphology and the Extradimensional Bypass

Josh C. Bongard

Chandana Paul

Artificial Intelligence Laboratory
University of Zürich
Zürich, Switzerland
{bongard, chandana}@ifi.unizh.ch

Abstract. In this paper, locomotion of a biped robot operating in a physics-based virtual environment is evolved using a genetic algorithm, in which some of the morphological and control parameters of the system are under evolutionary control. It is shown that stable walking is achieved through coupled optimization of both the controller and the mass ratios and mass distributions of the biped. It was found that although the size of the search space is larger in the case of coupled evolution of morphology and control, these evolutionary runs outperform other runs in which only the biped controller is evolved. We argue that this performance increase is attributable to extradimensional bypasses: adaptive ridges in the fitness landscape, created by the evolutionary control over morphology, that connect otherwise separated, sub-optimal adaptive peaks. In a similar study, a different set of morphological parameters are included in the evolutionary process. In this case, no significant improvement is gained by coupled evolution. These results show that the inclusion of the correct set of morphological parameters improves the evolution of adaptive behaviour in simulated agents.

1 Introduction

In the field of robotics, much work has been done on optimizing controllers for biped robots [1, 11, 21]. Similarly, genetic programming [8] and genetic algorithms [7] have been used to evolve controllers for hexapod robots. Genetic algorithms have also been used to evolve recurrent neural networks for bipedal locomotion: Fukuda et al [6] employed a dynamic simulator; Reil and Husband [18] employed a three-dimensional physics-based simulator. However, in all of these approaches, little or no consideration was paid to the mechanical construction of the agent or robot.

Alternatively, Brooks and Stein [3] and Pfeifer and Scheier [17] have pointed to the strong interdependence between the morphology and control of an embodied agent: design decisions regarding either aspect of an agent strongly bias the resulting behaviour. One implication of this interdependence is that often, a good choice of morphology can lead to a reduction in the size or complexity of the controller. For example, Lichtensteiger and Eggenberger [12] demonstrated that an evolutionary algorithm can optimize the sensor distribution of a mobile

robot for certain tasks, while the controller remains fixed. As an extreme case, the study of passive dynamics has made clear that a careful choice of morphology can lead to locomotion without any actuation or controller at all [16].

Examples now abound that demonstrate the evolution of both the morphology and control of simulated agents [19, 20, 10, 4, 15], as well as real-world robots [14, 9, 13] is possible. However, we argue in [2] that the coupled evolution of both morphology and control of adaptive agents is not as interesting in and of itself, but rather the implications of such studies open up a host of research questions regarding the evolution of adaptive behaviour that are not amenable to study solely through the optimization of control. Virtual Embodied Evolution (VEE) was introduced as a systematic methodology for investigating the implications of evolving both the morphology and control of embodied agents. In this paper we show not only that coupled evolution of both morphological and control parameters of a bipedal agent can facilitate the discovery of stable locomotion—despite the increased size of the search space necessitated by the inclusion of the additional morphological parameters—but also that only certain sets of morphological parameters facilitate evolutionary search.

The following section introduces the mechanical construction and neural controller of the biped agent, as well as the genetic algorithm used to evolve locomotion. Section 3 presents the results obtained from evolving only the neural networks for a bipedal agent, as well as evolutionary runs in which morphological parameters were included in the genome. Section 4 provides some discussion and analysis as to why coupled evolution of morphology and control can outperform the evolution of control. In the final section we conclude by stressing the importance of incorporating morphological considerations into the evolutionary investigation of adaptive behaviour.

2 The Model

For all of the evolutionary runs reported in this paper, the agents act within a physically-realistic, three-dimensional virtual environment¹. The agent is a simulation of a five-link biped robot with six degrees of freedom. The agent has a waist, and two upper and lower leg links as shown in Fig. 1 a. 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. In the second set of experiments reported in section 3, a second type of biped is used, in which five mass blocks are attached to the lower legs, upper legs and waist as shown in Fig. 1 b.

The joints are limited in their motion using 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 $-\frac{\pi}{7}$ and $\frac{\pi}{7}$ radians with respect to the vertical. The hip

¹ The environment and biped agents were constructed and evaluated using the real-time physics-based simulation package produced by MathEngine PLC, Oxford, UK, www.mathengine.com.




Fig. 1. Agent construction and neural network topology. a) shows the biped agent without the attached masses. b) shows the agent with the attached masses. c) gives a pictorial representation 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.

pitch joint has a range of motion between $-\frac{\pi}{10}$ and $\frac{\pi}{10}$, also with respect to the vertical. The knee joint has a range of motion between $-\frac{\pi}{2}$ and 0 with respect to the axis of the upper leg link to which it is attached. Table 1 summarizes the morphological parameters for both types of bipeds.

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 a recurrent neural network; 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 and recurrently connected, plus an additional bias neuron. The hidden and bias neurons are fully connected to the eight neurons in the output layer. Neuron activations are scaled by the threshold function $\frac{2}{1+e^{-a}} - 1$. 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.

Evolution of bipedal locomotion is achieved using a floating-point, fixed-length genetic algorithm. Each genome encodes weights for the 60 synapses composing the neural network, plus any additional morphological parameters. All values in the genome range between -1.00 and 1.00 . Each evolutionary run reported in this section is performed using a population size of 300, and is run

Table 1. The default size dimensions, masses and joint limits of the biped. All lengths and masses of the biped are relational: the unit length (**ul**), and the default mass (**um**), are set to the radii and masses of the knees and hip sockets, respectively. Parameters set in boldface indicate those parameters that are modified by evolution in the experiments reported in section 3. The valid ranges for these parameters are also given.

Index	Object	Dimensions	Mass
1	Knees	$r = 1\text{ul}$	1um each
2	Hip sockets	$r = 1\text{ul}$	1um each
3	Feet	$r = 2\text{ul}, w = 3\text{ul}$	1um each
4	Lower Legs	$r = [0.2, 0.8] \text{ ul}, h = 8\text{ul}$	0.25um each
5	Upper Legs	$r = [0.2, 0.8] \text{ ul}, h = 8\text{ul}$	0.25um each
6	Waist	$r = [0.2, 0.8] \text{ ul}, w = 8\text{ul}$	0.25um
7	Waist Block	$l = [0.4, 3.6] \text{ ul}, w = h = [0.2, 3.0] \text{ ul}$	0.103um
8	Lower Blocks	$l = [0.4, 3.6] \text{ ul}, w = h = [0.2, 3.0] \text{ ul}$	0.103um each
9	Upper Blocks	$l = w = [0.2, 3.0] \text{ ul}, h = [0.4, 3.6] \text{ ul}$	0.103um each

Index	Joint	Plane of Rotation	Range (rads)
10	Knee	sagittal	$-\frac{\pi}{2} \rightarrow 0$
11	Hip	sagittal	$-\frac{\pi}{7} \rightarrow \frac{\pi}{7}$
12	Hip	frontal	$-\frac{\pi}{10} \rightarrow \frac{\pi}{10}$

for 300 generations. Strong elitism is employed in which 150 of the most fit genotypes are preserved into the next generation. Tournament selection, with a tournament size of three, is employed to select genotypes from among this group 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: an average of five point mutations are introduced into each of these new genotypes, using random replacement.

In the set of experiments using the agent shown in Fig. 1 a, three additional morphological parameters are included in the genome. These parameters dictate the radii of the lower legs, upper legs and waist, respectively. The range of possible radii for these segments is $[0.2, 0.8]\text{ul}$. In the second set of experiments, eight morphological parameters are included in the genome. The first three values dictate the widths of the lower mass block pair, upper mass block pair and waist mass block, respectively, each of which can range between 0.2 and 3.0 ul. The next three values indicate the lengths of the lower mass block pair, upper mass block pair and waist mass block, respectively, which range between 0.4 and 3.6 ul. The final two values indicate the vertical placement of the two block mass pairs, which can range between 0.8 to 7.2 ul above the centre of the foot: the horizontal position of the waist block mass remains centred, and is not changed. In this way, all four blocks can be attached to the upper or lower pairs of legs. In the case of agents without block masses, the morphological parameter settings can affect the total mass, mass distribution and moment of inertia of the agent. In the case of agents with block masses, the morphological parameter values can affect only the mass distribution and the moment of inertia, although more

Table 2. Experimental regime summary.

Run Set	Morphology	Blocks	Total block mass	Genome length	Number of independent runs
1	Fixed	Absent	N/A	60	30
2	Variable	Absent	N/A	63	30
3	Fixed	Present	0.512um	60	20
4	Variable	Present	0.512um	68	20

degrees of freedom of the rotational moment of inertia are subjected to selection pressure in this case. For the variable morphology evolutionary runs, the three or eight morphological parameters are distributed evenly across the length of the genome in order to maximize recombination of these values during crossover.

The fitness of a genome is determined as follows. The weights encoded in the genotype are assigned to the synapses in the neural network, and in the case of the variable morphology bipeds without mass blocks, the radii of the waist, lower and upper legs are set based on the additional three values in the genome. In the case of the variable morphology bipeds with the mass blocks, the dimensions and positions of the blocks are set based on the additional eight parameters. The agent is then evaluated for up to 2000 time steps in the physical simulator. Evaluation halts prematurely if both of the feet leave the ground at the same time (this discourages the evolution of running gaits); the height of the waist passes below the height of the knees; or the waist twists more than 90 degrees away from the desired direction of travel. The northern distance of the agent at the termination of the evaluation period is then treated as the fitness of the genome.

3 Results

Four sets of evolutionary runs were conducted using the parameters given in Table 2. Fig. 2 summarizes the evolutionary performance of the two sets of runs using agents without mass blocks, and Fig. 3 reports the evolutionary performance of the two sets of runs using agent populations with mass blocks. It can be seen in Figs. 2 a and b that in both fixed and variable morphology agent populations, there is a roughly uniform distribution of fitness performance achieved by the most fit agents at the end of the runs. However Fig. 2 b indicates that variable morphology populations repeatedly achieved higher fitness values than the fixed morphology populations. Similarly, Figs. 2 c and d indicate that more fixed morphology populations do not realize any fitness improvement over evolutionary time, compared with the variable morphology populations.

In contrast, Fig. 3 indicates that stable locomotion is more difficult for evolution to discover for agent populations with fixed mass blocks, compared to agent populations without mass blocks, irrespective of whether or not the size and position of the blocks is under evolutionary control. Only two of the 20 populations achieve stable locomotion in both cases; the remaining runs do not realize any significant fitness improvements over evolutionary time.

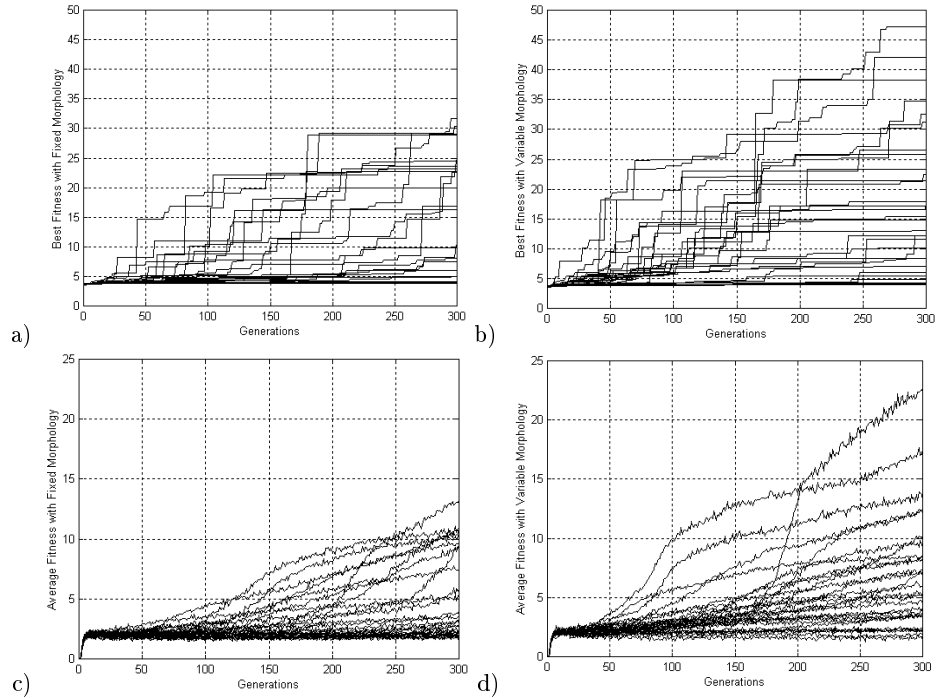


Fig. 2. Evolutionary performance of fixed and variable morphology agent populations without mass blocks. a) and b) report the highest fitness values attained by agents with fixed and variable morphologies, respectively, from 30 independently evolving populations of each agent type. c) and d) report the average fitness of these populations.

Fig. 4 a presents the average evolutionary performance of all the evolving fixed and variable morphology agent populations without mass blocks. It is shown that, independent of the initial random population, variable morphology populations tend to outperform fixed morphology populations. Fig. 4 b indicates that for the case of agent populations with mass blocks, on average there is no improvement in evolutionary search in the variable morphology populations over the fixed morphology populations.

In one evolutionary run from experiment set 2, the most fit agent at generation 170, with a fitness of 15.03, was replaced at generation 171 as the most fit agent in the population by its child, which sustained eight point mutations, and achieved a fitness of 23.17. The trajectories of the centres of mass of these two agents are indicated in Fig. 5 by the light gray and dark gray lines, respectively. Of the eight mutations, one of these was a morphological change that increased the radii of the child's lower leg pair from 0.578ul to 0.8ul. A third agent was tested, which was genotypically equivalent to the more fit child, except that the morphological mutation was suppressed. This third agent achieved a fitness of 20.87, and the trajectory of its centre of mass is indicated by the black line in Fig. 5.

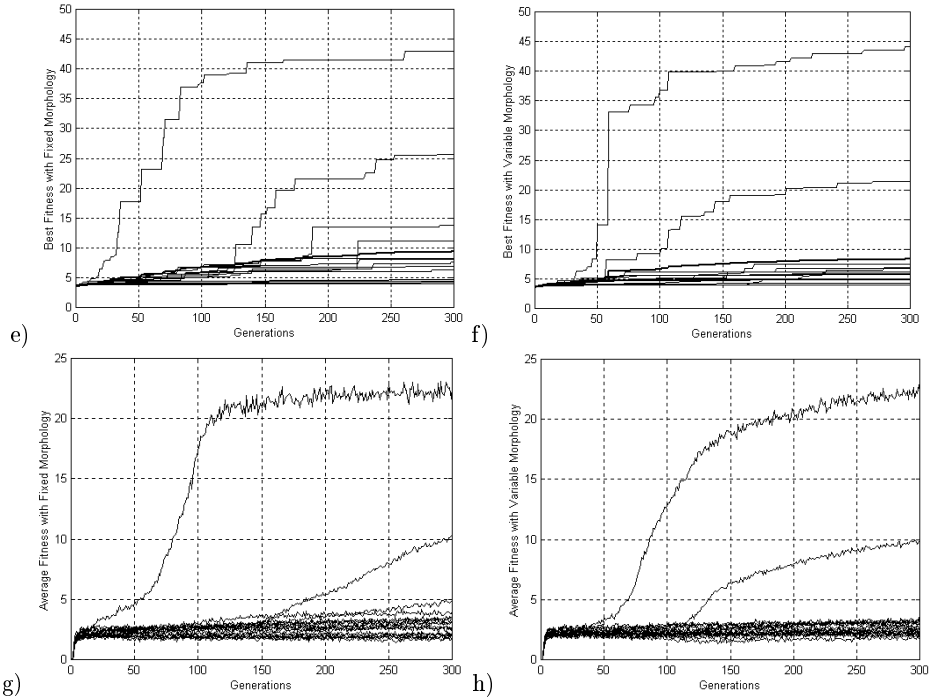


Fig. 3. Evolutionary performance of fixed and variable morphology agent populations with mass blocks. a) and b) report the highest fitness values attained by agents with fixed and variable mass blocks, taken from 30 independent evolutionary runs. c) and d) report the average fitness values of these populations.

4 Discussion

It is clear from Figs. 2 and 4 a that agent populations with varying leg widths tend to outperform agent populations with fixed leg widths. This stands in contrast to the intuitive notion that in the variable morphology case, the increased dimensionality of the search space—corresponding to the additional three morphological parameters—will degrade search. Furthermore, it is to be noted that the magnitude of morphological change possible through modification of these parameters is quite small: the minimum vertical centre of mass of an agent, relative to body height, is 0.49; the maximum vertical centre of mass is 0.59.

However, as made clear in Fig. 5, morphological mutations can have dramatic effects on an agent’s performance. In that case, the combination of seven control mutations and one morphological mutation realize a fitness increase of $23.17 - 15.03 = 8.14$, whereas only the control mutations give a fitness increase of $20.87 - 15.03 = 5.84$, resulting in the morphological mutation providing a fitness contribution of $23.17 - 20.87 = 2.3$. As can be seen from the trajectories of these agents, the control changes helped to stabilize the oscillations of the agent’s gait, as well as correct its direction of motion; the morphological mutation helped to further correct the direction of motion.




Fig. 4. Average evolutionary performance of fixed and variable morphology agents. a) indicates the average fitness of the fixed morphology and variable morphology populations reported in Fig. 2. b) indicates the average fitness of the populations reported in Fig. 3.




Fig. 5. Trajectories for three agents. The light gray line indicates the trajectory of the centre of mass of an agent that achieves a fitness of 15.03. This agent was succeeded as the most fit agent in the population by its child, which sustained eight point mutations: the trajectory of the centre of mass of this agent is indicated by the dark gray line. A third agent was tested, which was genotypically equivalent to the more fit child, except that the single morphological mutation was suppressed: the trajectory of this agent is indicated by the black line.

We did not find evidence that the variable morphology populations tended to converge on mass distributions far removed from the fixed setting, nor do they consistently converge on the same mass distribution. The vertical centres of mass of the most fit agents at the end of each run range between 0.52 and 0.57, and fall within this range with a roughly uniform distribution. This suggests that for our particular instantiation of bipedal locomotion and choice of controller, no one mass distribution is better than another. In other words, evolution of variable morphology agents does not perform better because evolution is able to discover a “good” morphology: rather, the addition of morphological parameters transforms the topology of the search space through which the evolving population moves, creating connections in the higher dimensional space between

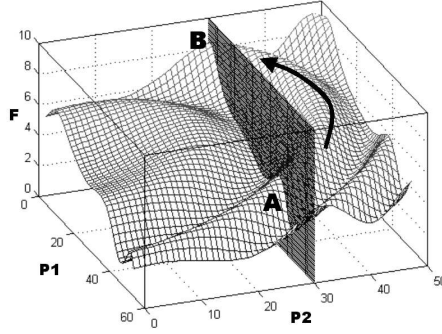


Fig. 6. Schematic representation of an extradimensional bypass. In the one-dimensional fitness landscape indicated by the cross-section within the vertical plane, the adaptive peak *A* is separated by a wide gulf of low fitness phenotypes from the higher peak *B*. In the higher dimensional fitness landscape indicated by the surface, an extradimensional bypass, represented by the curved surface, connects peaks *A* and *B*.

separated adaptive peaks in the lower dimensional space. These connections are known as extradimensional bypasses, and were introduced by Conrad in [5].

The cross-section within the vertical plane in Fig. 6 indicates a one-dimensional landscape in which the value of a single phenotypic trait *P1* dictates fitness *F*. This landscape contains two separated adaptive peaks, *A* and *B*: a population centred around peak *A* cannot easily make the transition to the higher fitness peak at *B*. However, through the addition of a second phenotypic parameter *P2*, the landscape is expanded to two dimensions (indicated by the surface), and an adaptive ridge—indicated by the upward sloping arrow—provides an opportunity for an evolving population to move from peak *A* to *B* via this extradimensional bypass.

We hypothesize that although the additional morphological parameters increase the dimensionality of the search space, in this case they introduce more adaptive ridges between local adaptive peaks, thereby smoothing the fitness landscape and facilitating evolutionary search. In other words, given a particular morphology, any combination of control changes does not confer increased fitness, but a change in morphology, coupled or followed by control changes does confer increased fitness. This is supported by the variable morphology populations, which do not converge on morphologies far removed from the default case. The hypothesis is also supported by the mutational event depicted in Fig. 5, in which higher fitness is achieved by modifications to both control and morphology.

However, the evolving agent populations with affixed mass blocks, indicated in Figs. 3 and Fig. 4 b, present a much different picture. In these populations, the addition of eight morphological parameters does not improve evolutionary search. In the 20 fixed morphology populations and 20 variable morphology populations, only two instances of stable locomotion were discovered in each. It is clear that bipedal locomotion using agents with mass blocks, using our

experimental set-up, is a more difficult task for the genetic algorithm, but the appearance of stable walking indicates it is not impossible for either the fixed or variable morphology regime to discover stable locomotion.

From our current experiments it is not clear why evolutionary search is not improved in this case, but it seems likely that there are two factors hindering improvement in the variable morphology populations. First, it seems plausible that the ruggedness of the lower dimensional fitness landscape, in the case of agents with fixed blocks, is greater than in the landscape for agents without mass blocks and fixed leg widths, because of the decreased evolutionary performance shown in Figs. 3 a and c, compared with the performance shown in Figs. 2 a and c. Second, the dimensionality of the search space for agent populations with mass blocks increases from 60 to 68, as compared with an increase of only 60 to 63 for agent populations without mass blocks. Thus in the case of the search space for agent populations with variable mass blocks, more smoothing is required to compensate for the greatly enlarged space, and the high ruggedness of the original space.

5 Conclusions and Future Research Directions

In this paper, stable locomotion was evolved in embodied, bipedal agents acting within a three-dimensional, physically-realistic virtual environment. It has been demonstrated that, for the case of locomotion in these agents, the subjugation of certain morphological parameters to evolutionary search increases the efficacy of the search process itself, despite the increased size of the search space.

Preliminary evidence was provided which suggests that artificial evolution does not do better in the case of the variable morphology populations because it is able to discover better morphologies than those imposed in the fixed morphology populations, but rather because the type of parameters included in the search create adaptive ridges linking previously separate adaptive peaks.

However, a control set of experiments was provided in which a different set of morphological parameters were included in the genomes of the evolving populations. In these experiments, there was no performance increase in the search ability of the genetic algorithm. This suggests that for the artificial evolution of adaptive behaviour, the arbitrary inclusion of morphological parameters does not always yield better results.

In future studies we plan to investigate in more detail how the inclusion of morphological parameters transforms the fitness landscape of the evolving populations. Moreover, we hope to formulate a systematic method for predicting which morphological parameters of embodied agents can augment the evolutionary discovery of adaptive behaviour.

References

1. Benbrahim, H., Franklin, J. A.: Biped Dynamic Walking Using Reinforcement Learning. In: *Robotics and Autonomous Systems* **22**: (1997) 283–302.

2. Bongard, J. C., Paul, C.: Investigating Morphological Symmetry and Locomotive Efficiency using Virtual Embodied Evolution. In: , J.-A. Meyer et al (eds.), *From Animals to Animats: The Sixth International Conference on the Simulation of Adaptive Behaviour*: (2000) 420–429.
3. Brooks, R. A., Stein, L. A.: Building Brains for Bodies. In: *Autonomous Robots 1:1* (1994) 7–25.
4. Chocron, O., Bidaud, P. Evolving Walking Robots for Global Task Based Design. In: *Proceedings of the IEEE Congress on Evolutionary Computation*. (1999) 405–412.
5. Conrad, M.: The Geometry of Evolution. In: *Biosystems* **24**: (1990) 61–81.
6. Fukuda, T., Komota, Y., Arakawa, T.: Stabilization Control Of Biped Locomotion Robot Based Learning With GAs Having Self-Adaptive Mutation And Recurrent Neural Networks. In: *Proceedings of 1997 IEEE International Conference on Robotics and Automation*. (1997) 217–222.
7. Gallagher, J.C., Beer, R.D, Espenschied, K.S., Quinn, R.D.: Application of Evolved Locomotion Controllers to a Hexapod Robot. In: *Robotics and Autonomous Systems* **19**: (1996) 95–103.
8. Gruau, F., Quatramaran, K.: Cellular Encoding for Interactive Evolutionary Robotics. Technical Report, University of Sussex, School of Cognitive Sciences, Brighton, UK. (1996)
9. Juárez-Guerrero, J., Muñoz-Gutiérrez, S., Mayol Cuevas, W. W.: Design of a Walking Machine Structure using Evolution Strategies. In: *IEEE Intl. Conf. on Systems, Man and Cybernetics*. (1998) 1427–1432.
10. Komosinski, M., Ulatowski, S.: Framsticks: Towards a Simulation of a Nature-Like World, Creatures and Evolution. In: *Proceedings of Fifth European Conference on Artificial Life*. Springer-Verlag, Berlin, (1998) 261–65.
11. Kun, A., Miller, W. T. III: Adaptive Dynamic Balance of a Biped Robot using Neural Networks. In: *Proceedings of the IEEE International Conference on Robotics and Automation*. (1996) 240–245.
12. Lichtensteiger, L., Eggenberger, P.: Evolving the Morphology of a Compound Eye on a Robot. In: *Proceedings of the Third European Workshop on Advanced Mobile Robots*. Piscataway, NJ. (1999) 127–34.
13. Lipson, H., Pollack, J. B.: Automatic Design and Manufacture of Artificial Life-forms. In: *Nature* **406** (2000) 974–78.
14. Lund, H. H., Hallam, J., Lee. W.-P.: Evolving Robot Morphology. In: *Proceedings of the IEEE Fourth International Conference on Evolutionary Computation*. IEEE Press (1997).
15. Mautner, C., Belew, R. K.: Evolving Robot Morphology and Control. In: Sugisaka, M. (ed.), *Proceedings of Artificial Life and Robotics 1999 (AROB99)*, Oita (1999).
16. McGeer, T.: Passive dynamic walking. In: *Int. J. Robotics Research* **9**:2 (1990) 62–82.
17. Pfeifer, R., Scheier, C.: *Understanding Intelligence*. MIT Press, Cambridge, MA (1999).
18. Reil, T., Husbands, P.: Evolution of Central Pattern Generators for Bipedal Walking in a Real-time Physics Environment. *IEEE Transactions on Evolutionary Computation*, submitted (2001).
19. Sims, K.: Evolving 3D Morphology and Behaviour by Competition. In: *Artificial Life IV*. MIT Press, Cambridge, MA (1994) pp. 28–39.
20. Ventrella, J.: Explorations of Morphology and Locomotion Behaviour in Animated Characters. In: *Artificial Life IV*. MIT Press, Cambridge, MA (1996) 436–441.
21. Vukobratovic et al.: *Biped Locomotion: Dynamics, Stability, Control and Applications*. Springer Verlag, Berlin (1990).