Multi-Objectivity for Brain-Behavior Evolution of a Physically-Embodied Organism

Jason Teo and Hussein A. Abbass
Artificial Life and Adaptive Robotics (A.L.A.R.) Lab,
School of Computer Science, University of New South Wales,
Australian Defence Force Academy Campus, Canberra, Australia.
{j.teo,h.abbass@adfa.edu.au}

Abstract

In this paper, we present a pareto multi–objective approach for evolving the behavior and brain (an artificial neural network (ANN)) of embodied artificial creatures. We will attempt to simultaneously minimize the network size while maximizing horizontal locomotion. A variety of network sizes and behaviors were generated by the pareto approach. The best networks exhibited a higher level of sensory-motor coordination and the creature was able to maintain the walking behavior under different environmental setups.

Introduction

Multi-objective optimization has been previously introduced for action selection in behavior-based robotics using conventional theory (Pirjanian 1998) and for the design of a robot arm using evolutionary theory (Coello Coello, Christiansen, & Aguirre 1998). In this paper, a multi-objective approach is investigated for evolving *artificial neural networks* (ANNs) that act as controllers for the legged locomotion of a 3-dimensional, artificial quadruped creature simulated in a physics-based environment. The Pareto-frontier Differential Evolution (PDE) algorithm is used to generate a pareto set of ANNs that trades-off between locomotion behavior and brain (*ie.* neural network) size. The concept of pareto defines a partial order dominance over the set of solutions. A solution X is said to be non–dominated (pareto) iff there is no other solution Y in the population where Y is better than X wrt all objectives.

The operational dynamics of the evolved creatures are analyzed to provide an insight into how a variety of controllers with different behaviors emerge from the evolution. A comparison between a set of pareto controllers showed that different brain sizes exhibit noticeably different locomotion behaviors. Although this may seem obvious, maintaining a variety of brain sizes and behaviors is mainly contributed by the optimization of conflicting objectives. We also found that a much higher level of sensory-motor coordination is present in the best evolved controller. Finally we investigated the effects of environmental, morphological and nervous system changes on the artificial creature's behavior. Similar to other studies, certain changes were found to be detrimental to the creature's walking-path. In contrast to other studies, however, the creature was able to maintain the walking behavior in a large majority of the experiments.

Situated and Embodied Evolution

The study of evolving physically situated and embodied artificial creatures has been a hotbed of research in recent years. The availability and maturation of commercial-off-the-shelf physics engines coupled with the dramatic increase of personal computing power have encouraged research into this intriguing field of artificial life (Taylor & Massey 2001). Since the pioneering work of Sims' (Sims 1994), there were notably few significant advancements in this field until very recently.

Research in this area generally falls into two categories: (1) the evolution of controllers for creatures with fixed (Arnold 1997; Bongard & Pfeifer 2002; Gritz & Hahn 1997; Harvey 1997; Ijspeert 2000; Reeve 1999) or parameterized morphologies (Lee, Hallam, & Lund 1996; Paul & Bongard 2001), and (2) the evolution of both the creatures' morphologies and controllers simultaneously (Bongard 2002; Hornby & Pollack 2001; Komosinski & Rotaru-Varga 2000; Lipson & Pollack 2000; Ray 2000; Sims 1994; Taylor & Massey 2001). Some work has also been carried out in evolving morphology alone (Eggenberger 1997) and evolving morphology with a fixed controller (Lichtensteiger & Eggenberger 1999). Related work using wheeled robots have also shown promising results in robustness and the ability to cope with changing environments by evolving plastic individuals that are able to adapt both through evolution and lifetime learning (Floreano & Urzelai 1998; 2000).

The emphasis of most of these studies have been on the role of genetic encodings and how different types of genotype-phenotype representations allow for greater evolvability (Bongard 2002; Hornby & Pollack 2001; Komosinski & Rotaru-Varga 2001). There have also been some investigations into the role of fitness functions and how they affect the direction of the evolutionary process (Floreano & Urzelai 2000; Komosinski & Rotaru-Varga 2000; Ray 2000). A very recent investigation explored how morphological complexity itself affects the emergence of more complex behavior in artificial creatures (Bongard & Pfeifer 2002). Considerably little has been said about the role of controllers in the artificial evolution of such creatures. It has been noted that the potential of designing more complex artificial systems through exploitation of sensory-motor coordination remains largely unexplored (Nolfi & Floreano 2002). As such, there is currently a lack of understanding of how the evolution of controllers affects the evolution of morphologies and behaviors in physically simulated creatures. It remains unclear what properties of an artificial creature's controller allow it to exhibit the desired behavior. A better understanding of controller complexity and the dynamics of evolving controllers should pave the way towards the emergence of more complex artificial creatures with more complex morphologies and behaviors.

Much work in evolutionary robotics have focused on evolving controllers for wheeled locomotion (Lee, Hallam, & Lund 1996; Floreano & Urzelai 1998; 2000; Nolfi & Floreano 1998). Less work have been conducted on evolving controllers for legged locomotion such as (Arnold 1997; Reeve 1999). Here we are attempting to evolve controllers that can generate walking behaviors for a four-legged creature.

Methods

The Physics Simulator

The simulation is carried out in a physically realistic environment which allows for rich dynamical interactions to occur between the creature and its environment. This in turn enables complex walking behaviors to emerge as the creature evolves the use of its sensors to control the actuators in its limbs through dynamical interactions with the environment. Furthermore, the accurate modelling of the simulation environment plays a crucial part in producing artificial creatures that move and behave realistically in 3D (Taylor & Massey 2001). A dynamic rather than kinematic approach is paramount in allowing for effective artificial evolution to occur. Physical properties such as forces, torques, inertia, friction, restitution and damping need to be incorporated into the artificial evolutionary system. To this end, the Vortex physics engine (Critical Mass Labs 2002) was employed to generate the physically realistic artificial creature (Figure 1) and its simulation environment.

The artificial creature is a basic quadruped with 4 short legs. Each leg consists of an upper limb connected to a lower limb via a hinge (one degree-of-freedom) joint and is in turn connected to the torso via another hinge joint. The mass of the torso is 1 and each of the limbs is 0.5. The torso has dimensions of 4 x 1 x 4 and each of the limbs has dimensions of 1 x 1 x 1. The hinge

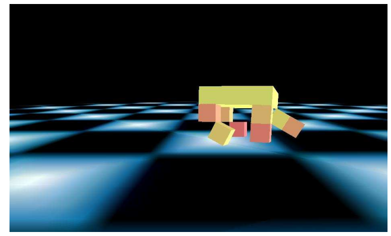

Figure 1: Screen capture of quadruped in the simulation environment.

joints are allowed to rotate between -1.57 to 0 radians for limbs that move counter-clockwise and 0 to 1.57 radians for limbs that move clockwise from their original starting positions. Each of the hinge joints are actuated by a motor that generates a torque producing rotation of the connected body parts about that hinge joint. The creature's overall central nervous system is illustrated in Figure 2.

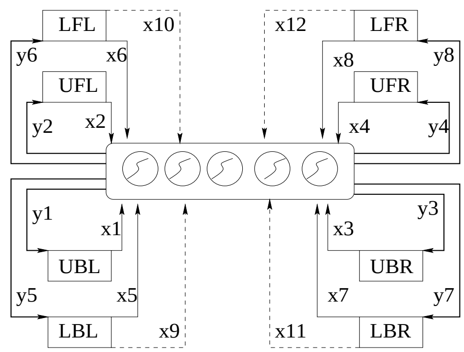

Figure 2: The quadruped's central nervous system.

Correspondingly, the artificial creature has 12 sensors and 8 actuators. The 12 sensors consist of 8 joint angle sensors $(x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8)$ corresponding to each of the hinge joints and 4 touch sensors $(x_9, x_{10}, x_{11}, x_{12})$ corresponding to each of the 4 lower limbs of each leg. The 8 actuators $(y_1, y_2, y_3, y_4, y_5, y_6, y_7, y_8)$ represent the motors that control each of the 8 articulated joints of the creature. These motors are controlled via outputs generated from the ANN controller which is then used to set the desired velocity of rotation of the connected body parts about that joint.

Controller Evolution Using PDE

Similar to (Abbass 2001; 2002), our chromosome is a class that contains one matrix Ω of real numbers representing the weights of the artificial neural network and one vector ρ of binary numbers (one value for each hidden unit) to indicate if a hidden unit exists in the network or not; that is, it works as a switch to turn a hidden unit on or off. The sum of all values in this vector represents the actual number of hidden units in a network. This representation allows simultaneous training of the weights in the network and selecting a subset of hidden units.

In the PDE algorithm for evolving ANNs, an entire set of controllers is generated in each evolutionary run without requiring any further modification of parameters by the user. The algorithm consists of the following steps:

1. Create a random initial population of potential solutions. The elements of the weight matrix Ω are assigned random values according to a Gaussian distribution $N(0,1)$. The elements of the binary vector ρ are assigned the value 1 with probability 0.5 based on a randomly generated number according to a uniform distribution between $[0,1]$; otherwise 0.

2. Repeat
   (a) Evaluate the individuals in the population and label those who are non-dominated.
   (b) If the number of non-dominated individuals is less than 3 repeat the following until the number of non-dominated individuals is greater than or equal to 3:
      i. Find a non-dominated solution among those who are not labelled.
      ii. Label the solution as non-dominated.
   (c) Delete all dominated solutions from the population.
   (d) Repeat
      i. Select at random an individual as the main parent α_1, and two individuals, α_2, α_3 as supporting parents.
      ii. **Crossover:** with some probability Uniform(0,1), do

$$\omega_{\text{ih}}^{\text{child}} \leftarrow \omega_{\text{ih}}^{\alpha_1} + N(0,1)(\omega_{\text{ih}}^{\alpha_2} - \omega_{\text{ih}}^{\alpha_3}) \tag{1}$$

$$\rho_{\text{h}}^{\text{child}} \leftarrow \begin{cases} 1 & \text{if } (\rho_{\text{h}}^{\alpha_1} + N(0,1)(\rho_{\text{h}}^{\alpha_2} - \rho_{\text{h}}^{\alpha_3})) \geq 0.5 \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

otherwise

$$\omega_{\text{ih}}^{\text{child}} \leftarrow \omega_{\text{ih}}^{\alpha_1} \tag{3}$$

$$\rho_{\text{h}}^{\text{child}} \leftarrow \rho_{\text{h}}^{\alpha_1} \tag{4}$$

and with some probability Uniform(0,1), do

$$\omega_{\text{ho}}^{\text{child}} \leftarrow \omega_{\text{ho}}^{\alpha_1} + N(0,1)(\omega_{\text{ho}}^{\alpha_2} - \omega_{\text{ho}}^{\alpha_3}) \tag{5}$$

otherwise

$$\omega_{\text{ho}}^{\text{child}} \leftarrow \omega_{\text{ho}}^{\alpha_1} \tag{6}$$

where each weight $(\omega_{\text{ih}}, \omega_{\text{ho}})$, and hidden unit flag, ρ, in the main parent are perturbed by adding to them a ratio, $F \in N(0,1)$, of the difference between the two values of this variable in the two supporting parents. At least one variable must be changed.

      iii. **Mutation:** with some probability Uniform(0,1), do

$$\omega_{\text{ih}}^{\text{child}} \leftarrow \omega_{\text{ih}}^{\text{child}} + N(0, \text{mutation\_rate}) \tag{7}$$

$$\omega_{\text{ho}}^{\text{child}} \leftarrow \omega_{\text{ho}}^{\text{child}} + N(0, \text{mutation\_rate}) \tag{8}$$

$$\rho_{\text{h}}^{\text{child}} \leftarrow \begin{cases} 1 & \text{if } \rho_{\text{h}}^{\text{child}} = 0 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

   (e) Until the population size is M

3. Until maximum number of generations is reached.

Experiments

Experimental Setup

A total of 480 evolutionary runs were conducted with varying population sizes, crossover rates, and mutation rates while fixing the fitness evaluation window to 500 timesteps. The crossover rates used were 0, 0.1, 0.2, 0.5 and 1 and the mutation rates used were also 0, 0.1, 0.2, 0.5 and 1. The evolutionary setup with a crossover rate of 0 and a mutation rate of 0 was omitted since this setup does not generate any variability at all in the population. The maximum number of hidden units permitted in evolving the artificial neural network was fixed at 15 nodes. Each experimental setup was repeated using 10 different seeds to allow the artificial evolution to commence from different starting points in the search space. Two population sizes of 20 and 30 were used with the corresponding number of generations being 30 and 20 respectively. The use of a small population size and number of generations is a feature of PDE since genetic diversity is naturally maintained by the pareto selection mechanism. The total number of genotypes over the entire span of the evolutionary process was kept constant at 600 genotypes in both these setups. In the final set of experiments, the creature was subjected to environmental, morphological and nervous system changes to observe the resultant change in its behavior. The details of these changes are presented along with the results and discussions of these experiments.

Results and Discussion

Evolutionary Parameters First we analyzed the effect of population size on the evolved locomotion behaviors. Overall, there did not appear to be any obvious differences in the range and quality of the evolved controllers between population sizes of 20 and 30. Both produced a considerably similar quality of locomotion behaviors although a larger population size did seem to produce controllers that were slightly better in terms of average locomotion fitness. However, the difference was not significant to investigate larger populations.

Different combinations of crossover and mutation rates did appear to produce results that varied across two broad spectrums. With both population sizes of 20 and 30, two distinct groups of controllers were generated through the evolutionary process: (1) runs that produced high quality solutions but with a low spread of genotypes, and (2) runs that produced mediocre solutions with a high spread of genotypes. Again, the quality of solutions refers to the average locomotion fitness and the spread of genotypes refers to the number of ANNs with different sizes in terms of hidden units. The first group of pareto optimal solutions with high quality and low spread were observed when fairly low mutation rates of 0.1 and 0.2 were used in combination with a low to medium crossover rate of between 0.1 to 0.5. The second group of pareto optimal solutions with lower quality but with a much wider spread of controller sizes were observed when a high mutation rate of 1 was used.

Operational Dynamics In this section, we analyze five pareto optimal controllers resultant from a typical run. To conduct this analysis, the ANNs were used individually to control the quadruped and the simulation period was extended to 5000 timesteps. This enables analysis of not only the evolved behavior but also its behavior beyond the fitness evaluation window.

The correlation analysis of the best evolved controller with 4 hidden units has 6 strongly positive correlation coefficients (>0.7). This indicates that the creature has evolved an ANN that has learned how to coordinate the movement of 7 sets of its limbs in order to achieve the most successful locomotion behavior among the pareto optimal controllers. With a correlation of 0.98, there is almost perfect coordination between the upper front left (UFL) and lower front right (LFR) limbs. Another almost perfectly coordinated motion comes from the upper back left (UBL) and upper back right (UBR) limbs with a correlation of 0.95. There is also a high level of correlation between the upper front left (UFL) and upper front right (UFR), lower front left (LFL) and upper front left (UFL), lower front left (LFL) and upper front right (UFR), upper front right (UFR) and lower front right (LFR), and lower front right (LFR) and lower front left (LFL) limbs.

Figure 3 illustrates the correlation between the 8 limbs during motion over 5000 timesteps along with the number of times each leg makes contact with the ground. Negative and positive correlation coefficients are drawn in dashed and solid lines respectively.

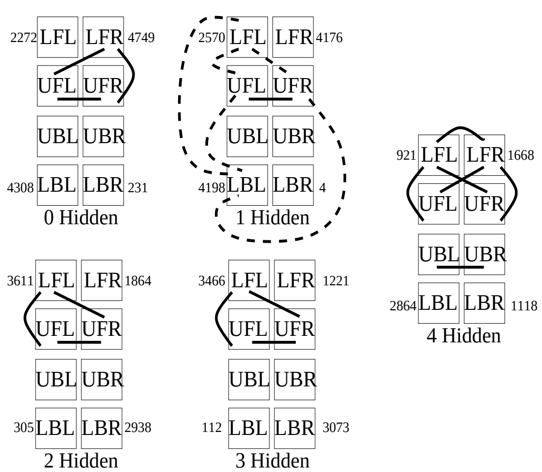

Figure 3: Illustration of correlation between limbs for pareto optimal controllers.

Analysis of the less successful pareto optimal networks reveals that there is far less coordination achieved by these controllers. At most 3 strongly correlated sets of limb movements were obtained using these controllers compared to 7 strongly correlated sets of limb movements using the best evolved controller. Furthermore, 5 strongly negative correlations (<-0.8) were detected in the controller with 1 hidden unit. These limbs are not only uncoordinated but are generating forces that act in direct opposition to each other, thereby further hindering the creature's ability to move.

Finally, we analyze the path of movement that was taken by the creature in attempting to maximize its horizontal distance covered during the extended simulation window of 5000 timesteps. Here we compared the paths of all networks on the pareto-frontier of the last generation of controller evolution.

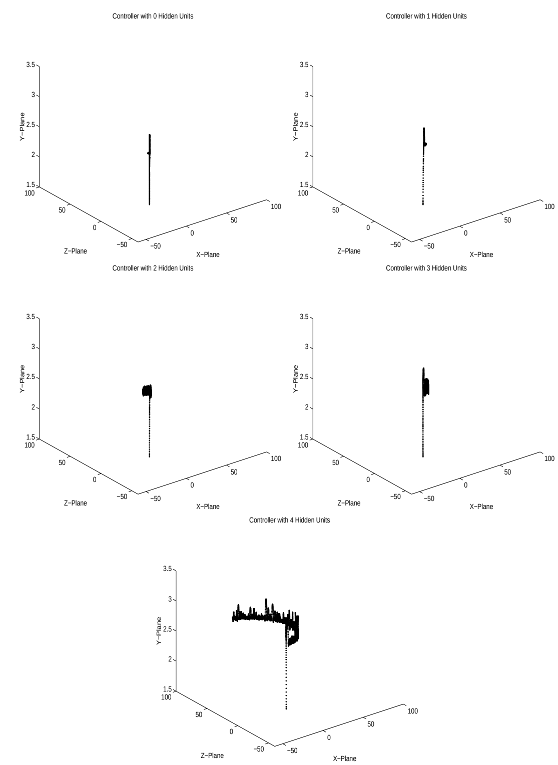

Figure 4: Path of movement using controller with 1. top left: 0; 2. top right: 1; 3. middle left: 2; 4. middle right: 3; and 5. bottom: 4; hidden units.

An interesting outcome from these multi-objective evolutions is that we get a range of controllers that vary in architectural complexity and locomotion capability. On the one hand, we have a totally random ANN with no hidden nodes (Figure 4.1) but which is still able to move the creature away from its origin, although the movement achieved within the stipulated 500 timesteps is extremely minimal (approximately 0.5m). In this random network, there is still an act of force on the creature permitting the small initial movement but it is unable to perform further locomotion due to the lack of synchronization ability. On the other hand, we have the best ANN that uses 4 hidden nodes (Figure 4.5) and is able to move almost 10m within the same time period. In addition, we have a further 3 ANNs (Figure 4.2,3,4) that utilize between 1 and 3 hidden nodes which again have differing locomotion capabilities. Thus, the multi-objective approach is able to provide the experimenter with a whole range of controllers within a single run that trades off between the individual optimization goals. This represents a significant advantage over single-objective evolutionary systems that need to be rerun multiple times in order to test the effect of other factors such as number of hidden units on the performance of artificial creatures (Bongard & Pfeifer 2002).

Effects of Friction In this section, we analyze the effects of changing some of the environmental parameters of the creature's world and observe the change in its behavior. Here, the same controller, which is the best evolved ANN with 4 hidden units, is used to control the creature across all different environmental conditions. The resultant behavior is again monitored over 5000 timesteps. First, we discuss the results obtained from changing the original frictional coefficient of 20 to lower values of 0, 5, 10 and 15. The purpose of this analysis was to investigate how the creature's ability to move would be affected by reduced amounts of grip with its locomotion surface.

The creature was not able to move horizontally at all with no ground friction. Its main movement here was mainly along the vertical direction as it attempted to stand up and repeatedly failed due to the lack of friction. With a very small friction of 5, the creature was able to move forwards although the overall distance travelled was less than in the original environment that had a significantly higher friction of 20. However, the path travelled in the environment with a friction of 5 was much straighter than in the original environment. This occurrence suggests that friction plays a larger role in making the creature turn compared to making it move forwards. From the next two environments which had increasingly higher frictions of 10 and 15, the overall trajectory of the paths begin to have more curvature as well as increasing overall distance travelled. Hence, it appears that varying locomotion surface conditions noticeably affect the creature's ability to walk both in terms of its trajectory as well as total distance travelled.

Effects of Gravity In this next section, we again change the environmental conditions but instead of surface condition, this time we change the world's gravitational field to approximately simulate conditions of that on the moon, Mars as well as Jupiter. The purpose of this set of experiments was again to see how the creature's behavior would be affected by environmental changes as well as exploring how hypothetical robots that are built under our planet's condition may be able to also function on numerous other planets that have significantly different gravities. Such robots may be desirable because firstly building them under normal terrestrial conditions will be significantly less complex than trying to simulate extra-terrestrial conditions. Secondly, if robots were able to perform reasonably independent of gravitational changes, then only a single group of similar robots need to be designed which would be able to explore multitudes of moons and planets with different surface gravities.

The creature was still able to function under the moon's much smaller gravity although the overall distance travelled was less than on Earth. There was also noticeably more vertical movement during the creature's locomotion as would be expected because of the smaller gravity. Under Mars' gravity, the creature's familiar U-shaped path becomes visible again although the overall distance travelled is again less than that achieved on Earth. The creature was significantly less successful under Jupiter's much higher gravity where after standing up, it was only able to move a small distance forward. From this analysis, it can be seen that the creature was still able to function under very different gravitational forces although it's locomotion was less successful than under Earth's normal gravity.

Effects of Morphological Changes Next, we analyze the change in the creature's behavior when there is a change in its morphology. Again the best evolved controller with 4 hidden units was used to control the creature and allowed to move for 5000 timesteps. In these experiments, we doubled the mass in certain parts of the creature's morphology.

Very pronounced changes were observed in the creature's locomotion behavior as a result of doubling the masses of all of its front limbs (Figure 5.1) and all of its back limbs (Figure 5.2). The doubling of mass in its front legs resulted in a locomotion path that had a straighter heading compared to the path observed with the original uniform mass distribution. Conversely, the doubling of mass in its back legs resulted in an even more pronounced curved locomotion trajectory than the original U-shaped path, where in this case the creature almost completed a full circle back to its original starting position. These phenomena may be explained by the fact that the creature achieved its locomotion from the coordinated movement of front limbs and back limbs respectively. As such, mass redistribution affecting entire front and back sections of the creature's body can be expected to result in significant changes to its locomotion behavior. The doubling of the creature's torso mass seemed to cause the creature's movement to head more directly towards the Z axis after making its initial left turn (Figure 5.3). The effect of doubling the mass of the front left and back right legs did not appear to alter the creature's path significantly except reducing the

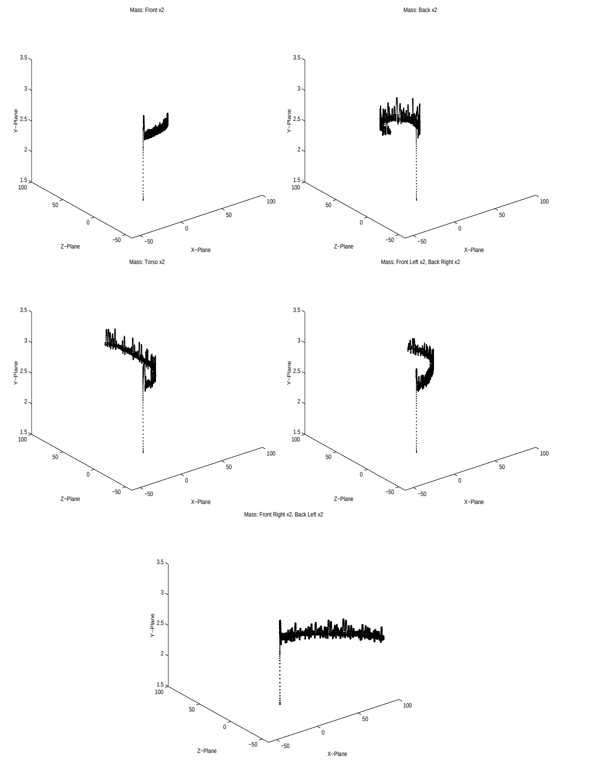

Figure 5: Path of movement with mass doubled in 1. top left: front legs; 2. top right: back legs; 3. middle left: torso; 4. middle right: front left and back right legs; 5. bottom: front right and back left legs.

magnitude and turning effect of its horizontal movement (Figure 5.4). The most pronounced change in the creature's overall heading was observed when the front right and back left legs were doubled in mass (Figure 5.5). This set of morphological changes appeared to have altered the nature of the creature's locomotion path from a predominantly left-turning trajectory to a right-turning trajectory. This may suggest that the contribution to overall movement from different legs are very different depending on the relative position of the legs with respect to the creature's body and direction of motion.

Effects of Sensory-Motor Failure In this section, we were interested in observing what would happen to the creature's locomotion behavior if some sensory–motor failure occurred. This would be akin to partial paralysis in four-legged animals where there is loss of sense and movement in some of their limbs. Here we disabled the joint angle and touch sensor as well as the hinge motors in the creature's entire front right limbs in the first setup and the entire back left limbs in the second setup. The best evolved controller with 4 hidden units was again used to operate the original creature with uniform mass distributions over 5000 timesteps.

Disabling the creature's front right leg seemed to have a harmful effect on its locomotion behavior. It struggled in trying to stand up and upon visual inspection of the simulation, this was explained by the fact it could not maintain its balance. As a result, the creature could not perform any horizontal movement at all. On the other hand, disabling the back left leg did not seem to cause as much harm to the creature's ability to move although its overall distance travelled was still significantly less compared to the original creature which had no impairments. In fact, upon closer inspection, the distinctive U-shaped locomotion pattern could still be observed but on a smaller scale. These analyzes again seem to suggest that the contribution of different legs to the overall locomotion behavior appeared to differ quite significantly depending on the position of the legs relative to the orientation of the creature's body and direction of movement. Thus disabling particular legs in certain positions resulted in dramatically different behaviors.

Conclusion

We have demonstrated the use of a multi-objective evolutionary algorithm for evolving artificial neural networks that act as controllers for the legged locomotion of an embodied and physically situated quadruped. We have shown that multi-objectivity allows for the natural maintenance of genetic diversity. The pareto-frontier that resulted from each single evolutionary run produced a set of ANNs that maximized the locomotion capabilities of the creature and at the same time minimized the size of the controller. The correlation and path analyzes of the pareto optimal controllers in operation provided an insight into how the complex coordination between the quadruped's different limbs generated the emergent locomotion behavior. Finally, we also observed that certain environmental, morphological and nervous system changes markedly affected the creature's overall locomotion behavior and in some cases caused total failure of its horizontal locomotion capability.

References

Abbass, H. A. 2001. A memetic pareto evolutionary approach to artificial neural networks. In Stumptner, M.; Corbett, D.; and Brooks, M., eds., *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence*, 1–12. Berlin: Springer-Verlag.

Abbass, H. A. 2002. An evolutionary artificial neural network approach for breast cancer diagnosis. *Accepted to appear in Artificial Intelligence in Medicine.*

Arnold, D. 1997. Evolution of legged locomotion. Unpublished masters thesis, School of Computing Science, Simon Fraser University, Burnaby, Canada.

Bongard, J. C., and Pfeifer, R. 2002. A method for isolating morphological effects on evolved behavior. In *Accepted to appear in Proceedings of the Seventh International Conference on the Simulation of Adaptive*

Behavior.

Bongard, J. C. 2002. Evolving modular genetic regulatory networks. In *Accepted to appear in Proceedings of the 2002 Congress on Evolutionary Computation.*

Coello Coello, C. A.; Christiansen, A. D.; and Aguirre, A. H. 1998. Using a new GA-based multiobjective optimization technique for the design of robot arms. *Robotica* 16:401–414.

Critical Mass Labs. 2002. Vortex. http://www.cm-labs.com.

Eggenberger, P. 1997. Evolving morphologies of simulated 3D organisms based on differential gene expression. In Husbands, P., and Harvey, I., eds., *Proceedings of the 4th European Conference on Artificial Life*, 205–213. Cambridge, MA: MIT Press.

Floreano, D., and Urzelai, J. 1998. Evolution and learning in autonomous mobile robots. In Mange, D., and Tomassini, M., eds., *Bio-Inspired Computing Machines: Towards Novel Computational Architectures.* Lausanne, Switzerland: PPUR, 1st edition. chapter 12, 317–364.

Floreano, D., and Urzelai, J. 2000. Evolutionary robotics: The next generation. In Gomi, T., ed., *Proceedings of Evolutionary Robotics III*, 231–266. Ontario: AAI Books.

Gritz, L., and Hahn, J. K. 1997. Genetic programming evolution of controllers for 3-D character animation. In Koza, J. et al. eds., *Genetic Programming 1997: Proceedings of the 2nd Annual Conference*, 139–146. San Francisco: Morgan Kauffman.

Harvey, I. 1997. Artificial evolution and real robots. *Artificial Life and Robotics* 1(1):35–38.

Hornby, G. S., and Pollack, J. B. 2001. Body-brain coevolution using L-systems as a generative encoding. In Spector, L. et al. eds., *Proceedings of the Genetic and Evolutionary Computation Conference*, 868–875. San Francisco: Morgan Kauffman.

Ijspeert, A. J. 2000. A 3-D biomechanical model of the salamander. In Heudin, J.-C., ed., *Proceedings of the 2nd International Conference on Virtual Worlds*, 225–234. Berlin: Springer-Verlag.

Kauffman, S. A. 1993. *The Origins of Order.* New York: Oxford University Press.

Komosinski, M., and Rotaru-Varga, A. 2000. From directed to open-ended evolution in a complex simulation model. In Bedau, M.; McCaskill, J.; Packard, N.; and Rasmussen, S., eds., *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, 293–299. Cambridge, MA: MIT Press.

Komosinski, M., and Rotaru-Varga, A. 2001. Comparison of different genotype encodings for simulated three-dimensional agents. *Artificial Life* 7(4):395–418.

Lee, W.-P.; Hallam, J.; and Lund, H. J. 1996. A hybrid GP/GA approach for co-evolving controllers and robot bodies to achieve fitness-specific tasks. In *Proceedings of the 3rd IEEE International Conference on Evolutionary Computation*, 384–389. Piscataway, NJ: IEEE Press.

Lichtensteiger, L., and Eggenberger, P. 1999. Evolving the morphology of a compound eye on a robot. In *Proceedings of the Third European Workshop on Advanced Mobile Robots*, 127–134. Piscataway, NJ: IEEE Press.

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

Nolfi, S., and Floreano, D. 1998. Learning and evolution. *Autonomous Robots* 7(1):89–113.

Nolfi, S., and Floreano, D. 2002. Synthesis of autonomous robots through evolution. *Trends in Cognitive Science* 6(1):31–36.

Paul, C., and Bongard, J. C. 2001. The road less travelled: Morphology in the optimization of biped robot locomotion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, 226–232. Piscataway, NJ: IEEE Press.

Pirjanian, P. 1998. Multiple objective action selection in behavior-based control. In *Proceedings of the 6th Symposium for Intelligent Robotic Systems.* Edinburgh, UK.

Ray, T. S. 2000. Aesthetically evolved virtual pets. In Maley, C., and Boudreau, E., eds., *Artificial Life 7 Workshop Proceedings*, 158–161. http://www.hip.atr.co.jp/∼ray/pubs/alife7a/.

Reeve, R. 1999. *Generating Walking Behaviors in Legged Robots.* Unpublished PhD thesis, University of Edinburgh, Scotland.

Sims, K. 1994. Evolving 3D morphology and behavior by competition. In Brooks, R., and Maes, P., eds., *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, 28–39. Cambridge, MA: MIT Press.

Taylor, T., and Massey, C. 2001. Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life* 7(1):77–87.