

COORDINATION AND SYNCHRONIZATION OF LOCOMOTION IN A VIRTUAL ROBOT

Jason Teo, 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

This paper investigates the use of a multi-objective approach for evolving artificial neural networks 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 optimal set of artificial neural networks that optimizes the conflicting objectives of maximizing locomotion behavior and minimizing neural network complexity. Here we provide an insight into how the controller generates the emergent walking behavior in the creature by analyzing the evolved artificial neural networks in operation. A comparison between pareto optimal controllers showed that ANNs with varying numbers of hidden units resulted in noticeably different locomotion behaviors. We also found that a much higher level of sensory-motor coordination was present in the best evolved controller.

Keywords: artificial evolution, artificial life, embodied cognitive science, evolutionary robotics.

1. INTRODUCTION

There has been a strong resurgence of research into the evolution of morphology and controller of physically simulated creatures. The pioneering and captivating work of Sims [17] in 1994 has not been paralleled until very recently. Further work in this area was limited by the complexity of programming a realistic physics-based environment and the steep computational resources required to run the artificial evolution. These physically realistic simulations of evolving artificial minds and bodies have become more accessible to the wider research community as a result of the recent convergence in the maturation of physics-based simulation packages and increase of raw computing power of personal computers [18].

Research in this area generally falls into two categories: (1) the evolution of controllers for creatures with fixed [4, 9] or parameterized morphologies [12, 16], and (2) the evolution of both the creatures' morphologies and controllers si-

multaneously [8, 11, 14, 18]. Some work has also been carried out in evolving morphology alone [6] and evolving morphology with a fixed controller [13]. Related work using mobile 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 [7].

However, 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 [15]. 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.

In this paper, we investigate the use of a multi-objective approach in evolving controllers for a fixed morphology artificial creature. By generating a pareto-frontier consisting of multiple ANNs with differing locomotion capabilities and varying architecture complexities, a comparison of controller size against behavior fitness can be made. This study will hopefully provide some insights into the architectural complexity of controllers required for generating walking behaviors in 3D, physically simulated creatures. A further advantage of using a multi-objective approach for artificial evolution is that genetic diversity is maintained naturally during the course of the evolutionary process. It has been observed that loss of genetic diversity causes problems in the artificial evolution of virtual creatures [10]. In this paper, the Pareto-frontier is used to evolve a pareto optimal set of artificial neural networks (ANNs) [1, 2] that act as controllers for the quadruped creature.

2. CONTROLLER EVOLUTION USING PDE

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 [18]. The Vortex physics engine [5] was employed to generate the physically realistic artificial creature and its simulation environment. A screen capture of the quadruped moving in its environment is shown in Figure 1.




Figure 1: Screen capture of quadruped in the 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 1kg and each of the limbs is 0.5kg. The torso has dimensions of 4 x 1 x 4m and each of the limbs has dimensions of 1 x 1 x 1m. The hinge 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.

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.

The Pareto-frontier Differential Evolution (PDE) algorithm [3] is used to evolve a pareto optimal set of artificial




Figure 2: The quadruped’s central nervous system. The three letter abbreviations identify each of the 8 different limbs. The first letter denotes (U)pper or (L)ower, the second denotes to (F)ront or (B)ack, and the third denotes (R)ight or (L)eft. Here only the value for the 4 lower limbs are shown as touch sensors are only located in these parts of the creature’s legs.

neural networks (ANNs) [1, 2] that act as controllers for the quadruped creature. An entire set of controllers is generated in each evolutionary run without requiring any further modification of parameters by the user. The artificial evolutionary system proceeds along two separate goals: to (1) maximize horizontal locomotion and, (2) minimize the complexity of the controller. In this initial study, controller complexity is measured using the number of hidden nodes that are used in the ANN. In future work, we intend to define more rigorous measures of controller complexity by taking into consideration other ANN architectural features such as number of connection weights as well as number of nodes in the input and output layers. The aim here is to produce a set of pareto optimal controllers that trades-off between locomotion capabilities and controller complexity.

Similar to [1, 2], 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.

3. EXPERIMENTS

3.1. 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 rate 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. The number of generations and population size were fixed at 20 and 30 respectively for the first set of runs. In the second set of runs, these parameter values were reversed to 30 for number of generations and 20 for population size to enable a fair comparison between the effect of the two population sizes (the total number of genotypes over the entire span of the evolutionary process was kept constant at 600 genotypes in both these setups).

3.2. Results and Discussion

High quality and a low spread of solutions were obtained with low mutation and low to medium crossover whereas mediocre solutions with a wider variety of controller sizes were obtained with high mutation and low to high crossover. Population size did not appear to significantly affect the quality and spread of pareto optimal solutions in these experiments although a very slight advantage in terms of quality and variety of controller sizes was observed with the larger population size of 30.

In the rest of this section, we analyze the 5 pareto optimal controllers in operation. To conduct these analyses, the best evolved ANNs described in the previous section 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. Table 1 lists the correlation coefficients between the joint angles of the respective limbs of the creature in motion over 5000 timesteps.

	UBL	UFL	UBR	UFR	LBL	LFL	LBR	LFR
UBL	1	-0.29	0.95	-0.11	-0.55	-0.29	0.09	-0.28
UFL		1	-0.24	0.73	-0.07	0.89	0.02	0.98
UBR			1	-0.07	-0.45	-0.24	0.09	-0.23
UFR				1	-0.13	0.88	0.02	0.71
LBL					1	-0.09	-0.04	-0.06
LFL						1	0.02	0.88
LBR							1	0.03
LFR								1

Table 1: Correlation coefficients between the joint angles of the creature’s 8 limbs in motion over 5000 timesteps with 4 hidden units. The three letters are as presented in Figure 2.

The correlation analysis of the best evolved controller with 4 hidden units has 7 strongly positive correlation co-

efficients (>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. In summary, the creature achieves locomotion by coordinating the movements between:

1. upper limbs of its back legs (0.95)
2. upper and lower limbs of its front left leg (0.89)
3. upper and lower limbs of its front right leg (0.71)
4. upper limbs of its front legs (0.73)
5. lower limbs of its front legs (0.88)
6. opposing limbs of its front legs (0.98, 0.88)

Some of these coordinated movements are quite obvious when inspecting the movement of the quadruped visually during simulation, for example the coordination present between the front legs and between the back legs. However, some coordinated movements are less obvious visually, for example the movements of opposing limbs in the front legs. Such complex coordinations are expected in locomotion of legged creatures, which largely explains why hand-designing controllers for such creatures tends to be extremely difficult and normally results in less than desirable behaviors. The illustrations that follow in Figure 3 graphically illustrate the correlation between the 8 limbs during motion over 5000 timesteps along with the number of times each leg makes contact with the ground.

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. It can be seen from the graphical illustration that the best evolved controller with 4 hidden units achieved high coordination between all of the creature’s front limbs as well as in one set of its back limbs. However, with all of the other less successful controllers, coordination was only achieved in some of its front limbs and no coordination was present at all in the back limbs. In these latter cases, the creature is only able to generate useful movements from

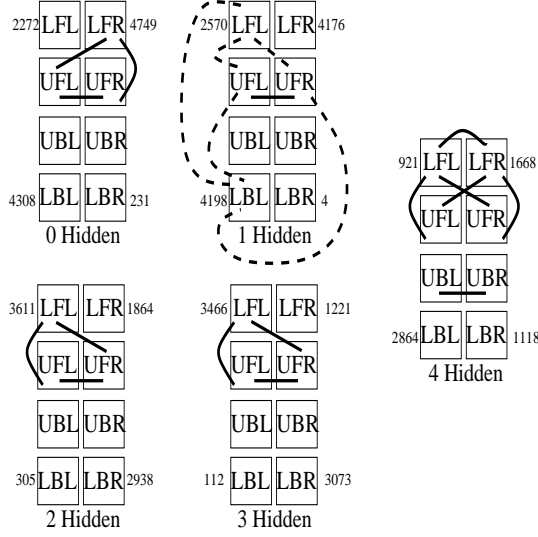


Figure 3: Illustration of correlation between limbs for pareto optimal controllers.

its front legs with no contribution at all from its back legs which resulted in poor locomotion behavior. 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.

Next, we analyze the synchronization between the touch sensors. The value used in this analysis represents the total number of times each pair of legs either contact the ground or is in the air, as explained in the equation below:

$$\bullet \text{ Touch} = \frac{\text{count}(x_i = x_j)}{\text{Total number of timesteps}}$$

	LBL	LFL	LBR	LFR
LBL	1	0.35	0.53	0.34
LFL		1	0.63	0.65
LBR			1	0.50
LFR				1

Table 2: Touch synchronization between the creature's legs in motion over 5000 timesteps with 4 hidden units. The three letter abbreviations identify each of the 8 different limbs.

The previous equation was used for all networks on the pareto frontier. The best spread of synchronization between pairs of legs are achieved in the controller with 4 hidden units, which demonstrated the best locomotion behavior, as shown in Table 2. This can be attributed to the fact that a balance between the number of times each leg synchronizes with a particular leg, for example to balance the body, as

well as with other legs, for example to push the creature forwards, needs to be achieved in order to generate useful locomotion. Looking at the controllers with less numbers of hidden units, a larger spread of synchronization can be noticed, which means that the creature has pairs of limbs that spend the majority of the time either balancing the body or attempting to push the creature forwards without striking a balance between these two critical aspects of successful locomotion.

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.

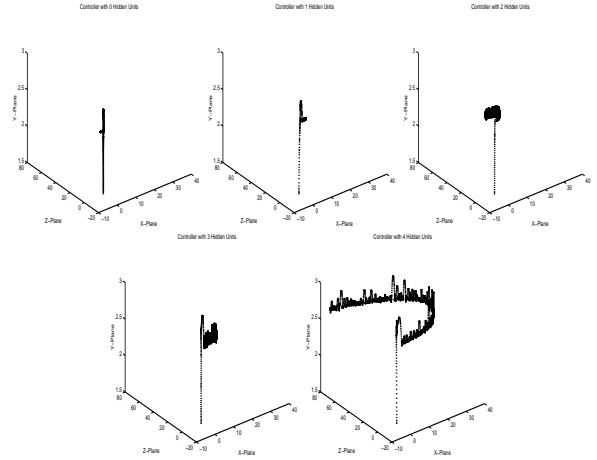


Figure 4: Path of movement using controller with 0, 1, 2, 3, and 4 hidden units.

As can be seen from the graphs depicting the movement of the creature, the least amount of movement was achieved by the controller with no hidden units. The creature was only able to partially stand up and hardly moved at all from its origin. Not much improvement was achieved by the controller that used 1 hidden unit. Its behavior was almost identical to the controller with no hidden units although it did manage to move slightly further away from its origin. We start to see significantly more movement with the controller with 2 hidden units where after standing up fairly efficiently, it manages to move in a small U-shape path away from its origin. Using the controller with 3 hidden units, the creature again manages to stand up very efficiently and follow a fairly straight path away from the origin. The distance covered using this controller was slightly more than the controller with 2 hidden units. Finally the best evolved controller which used 4 hidden units showed a significantly higher locomotion capability where it very successfully carved a large U-shape path along the X and Z planes starting from its origin. Using this controller, the creature first stood up very quickly and

moved in a reasonably straight line toward 10 m along the X plane during the first 500 timesteps, which represented the evaluation window during evolution. Beyond the evaluation window, the controller appears to veer the creature towards the Z plane and eventually turns around on its original path and heads in the reverse direction along the X plane. This shows that although the creature's controller performed well during the period where its fitness was subjected to evolutionary pressure, its long-term locomotion behavior beyond this point was noticeably different from the original intended behavior. Comparing across the controllers with different numbers of hidden units, we can also observe that controller complexity does in fact play a strong role in determining the emergent locomotion behaviors within the same creature. On one extreme, we have a controller with no hidden units that is only able to partially stand up and achieves virtually no horizontal movement to the other extreme where we have a controller with 4 hidden units that is able to not only stand up quickly but also move the creature over very large distances.

Another 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 but 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 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 and is able to move almost 10m within the same time period. In addition, we have a further 3 ANNs 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 re-run multiple times in order to test the effect of other factors such as number of hidden units on the performance of artificial creatures [4].

4. CONCLUSION

We have demonstrated a multi-objective approach to evolving artificial neural networks for controlling the locomotion of a 3D, physically simulated artificial creature. The pareto-frontier that resulted from each single evolutionary run provided a set of ANNs which maximized the locomotion capabilities of the creature and at the same time minimized the size of the controller. The correlation and path analyses 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. For future work, we intend to investigate the effects of controller complexity when both the morphology and controller are co-evolved simultaneously.

5. REFERENCES

- [1] Hussein A. Abbass. A memetic pareto evolutionary approach to artificial neural networks. In M. Stumptner, D. Corbett, and M. Brooks, editors, *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence (AI'01)*, pages 1–12, Berlin, 2001. Springer-Verlag.
- [2] Hussein A. Abbass. An evolutionary artificial neural network approach for breast cancer diagnosis. *Accepted to appear in Artificial Intelligence in Medicine*, 2002.
- [3] Hussein A. Abbass, Ruhul Sarker, and Charles Newton. PDE: A pareto-frontier differential evolution approach for multi-objective optimization problems. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC2001)*, volume 2, pages 971–978, Piscataway, NJ, 2001. IEEE Press.
- [4] Josh C. Bongard and Rolf Pfeifer. 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 (SAB2002)*, 2002.
- [5] Critical Mass Labs. Vortex. <http://www.cm-labs.com>, 2002.
- [6] Peter Eggenberger. Evolving morphologies of simulated 3D organisms based on differential gene expression. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life (ECAL97)*, pages 205–213, Cambridge, MA, 1997. MIT Press.
- [7] Dario Floreano and Joseba Urzelai. Evolutionary robotics: The next generation. In T. Gomi, editor, *Proceedings of Evolutionary Robotics III*, pages 231–266, Ontario, 2000. AAI Books.
- [8] Gregory S. Hornby and Jordan B. Pollack. Body-brain coevolution using L-systems as a generative encoding. In L. Spector, E. Goodman, A. Wu, W.B. Langdon, H.M Voight, M. Gen, S. Sen, M. Dorigo, S. Pezeshek, M. Garzon, and E. Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 868–875, San Francisco, 2001. Morgan Kauffman.

- [9] Auke J. Ijspeert. A 3-D biomechanical model of the salamander. In J.-C. Heudin, editor, *Proceedings of the 2nd International Conference on Virtual Worlds (VW 2000)*, pages 225–234, Berlin, 2000. Springer-Verlag.
- [10] Maciej Komosinski and Adam Rotaru-Varga. Comparison of different genotype encodings for simulated three-dimensional agents. *Artificial Life*, 7(4):395–418, 2001.
- [11] Maciej Komosinski and Szymon Ulatowski. Framsticks: Towards a simulation of a nature-like world, creatures and evolution. In D. Floreano, J.-D. Nicoud, and F. Mondada, editors, *Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life (ECAL99)*, pages 261–265, Berlin, 1999. Springer-Verlag.
- [12] Wei-Po Lee, John Hallam, and Henrik J. Lund. 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*, pages 384–389, Piscataway, NJ, 1996. IEEE Press.
- [13] Lukas Lichtensteiger and Peter Eggenberger. Evolving the morphology of a compound eye on a robot. In *Proceedings of the Third European Workshop on Advanced Mobile Robots (Eurobot '99)*, pages 127–134, Piscataway, NJ, 1999. IEEE Press.
- [14] Hod Lipson and Jordan B. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406:974–978, 2000.
- [15] Stefano Nolfi and Dario Floreano. Synthesis of autonomous robots through evolution. *Trends in Cognitive Science*, 6(1):31–36, 2002.
- [16] Chandana Paul and Josh C. Bongard. 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 (IROS2001)*, volume 1, pages 226–232, Piscataway, NJ, 2001. IEEE Press.
- [17] Karl Sims. Evolving 3D morphology and behavior by competition. In R. Brooks and P. Maes, editors, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 28–39, Cambridge, MA, 1994. MIT Press.
- [18] Tim Taylor and Colm Massey. Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life*, 7(1):77–87, 2001.