

Neuroevolution of Humanoids that Walk Further and Faster with Robust Gaits

Ben Jackson¹ and Alastair Channon¹

¹ School of Computing and Mathematics, Keele University, ST5 5BG, UK
{b.jackson,a.d.channon}@keele.ac.uk

Abstract

Bipedal locomotion requires precise rhythm and balance. Here we demonstrate two fitness-function enhancements applied to OpenAI’s 3D Humanoid-v1 walking task using a replica of Salimans *et al.*’s evolution strategy (Salimans *et al.*, 2017). The first enhancement reduces control cost, following a start-up period, and the second enhancement penalises poor balance. Individually, each enhancement results in improved gaits and doubles both median speed and median distance walked. Combining the two enhancements results in little further improvement in the absence of noise but is shown to produce gaits that are much more robust to noise in their actions, with median speed, distance and time two to five times those of the default and individual-enhancement gaits at an intermediate noise level.

Introduction

Bipedal locomotion has been the focus of many studies (Winter, 1991). Its emergence has been associated with endurance running and freeing the hands for other tasks such as tool use (Hewes, 1961; Bramble and Lieberman, 2004). The main challenges relate to balance and generating cyclic motion across two limbs to produce a stable gait. In robotics, bipedal walking is often achieved through zero-moment point computation (Vukobratović and Borovac, 2004), for example in evaluating fitness (Fukuda *et al.*, 1997).

Initially, the notable instances of bipedal learning utilised Central Pattern Generators. CPGs are recurrent neural networks that produce rhythmic activity (without requiring rhythmic inputs) that is typically modulated by descending or peripheral inputs (Guertin, 2013). Modelled after pattern generators found in human and animal spines, they offer a biologically inspired solution (Taga *et al.*, 1991; Ijspeert, 2008; Van der Noot *et al.*, 2015). Reil and Husbands used a genetic algorithm to evolve a CPG’s weights to produce stable bipedal walking (Reil and Husbands, 2002). Measuring fitness as distance from the point of origin, they achieved a 10% stability success rate. This was improved to 80% through the introduction of an oscillatory bonus in the fitness function. Reil and Husbands also added sensors to achieve directional walking. Directional motion was also exhibited

in (Gökçe and Akin, 2010), which used evolution strategies (Beyer and Arnold, 2001) to optimise CPG-based walking in simulation and then on real robots.

Following this, DeepMind used reinforcement learning to train a 3D humanoid model (Tassa *et al.*, 2012), using the MuJoCo physics engine (Emanuel Todorov *et al.*, 2012), to produce complex bipedal locomotion (Heess *et al.*, 2017). A distance-based reward function and policy-based gradient descent learning were used in incrementally complex environments to produce behaviours. The humanoid agents were able to demonstrate running, crouching, jumping and turning behaviours, the most notable successful model to date.

Salimans *et al.* (OpenAI) evolved walking gaits for the MuJoCo 3D Humanoid-v1 environment in OpenAI Gym (Tassa *et al.*, 2012; Brockman *et al.*, 2016) using an evolution strategy (Salimans *et al.*, 2017). The architecture described in their paper was a multilayer perceptron with two 64-unit hidden layers (using tanh units) mapping 376 inputs to 17 joint torques. In each iteration, many episodes were run using random parameter perturbations to test the robustness of the current strategy. The strategy parameters for the next iteration were then updated based on the calculated gradient estimate of the combined episode results. This produced successful biped neurocontrollers in as little as ten minutes, using 1440 CPU cores.

Subsequently, Petroski Such *et al.* (Uber AI Labs) used a genetic algorithm on the same problem (3D Humanoid-v1), using two 256-unit hidden layers (matching the configuration file included in the source code released by Salimans *et al.*), achieving success on this task but noting that their GA “took 15 times longer to perform slightly worse than ES” (Petroski Such *et al.*, 2018). They also noted that (while only just qualifying as a deep neural network, having more than one hidden layer) this network contains approximately 167k parameters, orders of magnitude greater than the previous largest neural networks evolved for robotics tasks. They encoded these parameters using a novel method that stores, for each genotype, an initialization seed and a list (that grows with each mutation) of random seeds used to generate mutations to the vector of parameters. They




Figure 1: The Humanoid-v1 walker used in this and previous work (Salimans et al., 2017; Petroski Such et al., 2018; Conti et al., 2018).

also applied a weighted novelty-search version of the ES, which was able to produce high scores in a deceptive environment that required the humanoid to walk around a trap (Conti et al., 2018) and demonstrated that ES is more robust to parameter perturbation in the humanoid locomotion task, than both their GA and Trust Region Policy Optimization (Lehman et al., 2018).

In this paper we report on two enhancements that enable Salimans *et al.*'s evolution strategy, using two 256-unit hidden layers, to produce populations of 3D Humanoid-v1 agents (figure 1) that walk further and faster and exhibit novel gaits with much increased robustness to the addition of noise to actions. We also report on the interaction between the two enhancements and how they overcome the difficulty of the task.

Methods

The methods featured here are applied to a replica of Salimans *et al.*'s evolution strategy, which was already capable of producing gaits that move 3D Humanoid-v1 agents quickly and efficiently enough to pass the humanoid walking task. The humanoid walker's goal is to travel (in any direction) as fast and efficiently as possible, failing when its torso falls below (or above) a certain height. Fitness is defined as the sum (over time) of four rewards/penalties that are computed at each timestep: a reward for linear velocity, a control cost based on energy expended, a cost based on how hard the humanoid impacts the ground, and a reward for standing. Our methods aim to improve average speed, distance travelled and episode duration (time). We also aim to improve robustness to noise in the application of actions. Each permutation of enhancements (and the default base-case) was run 20 times, each for 600 iterations (generations). The standard evolution strategy was used due to

its superior performance (see above). As mentioned above, the architecture maps 376 inputs (humanoid state variables: position, rotations, velocities, forces and inertia values) to 17 joint torques, from which we conclude approximately 30k parameters (weights and biases). Each set of 20 runs took around three days to evolve using 40 CPU cores. Each episode (walker fitness evaluation) was limited to a maximum of 1000 timesteps.

Fitness Function Enhancements

Fitness is the sum (over time) of four rewards/penalties that are computed at each timestep: a reward for linear velocity, a control cost based on energy expended, a cost based on how hard the humanoid impacts the ground, and a bonus for having not failed:

$$r = linvel - contcost - impcost + alivebonus$$

Even in the default case, without any enhancement, this led to the evolution of fast, efficient walking in line with previous results. The impact cost prevented hopping behaviours. However, a considerable number of evolved gaits involved a shuffling motion. These walkers slid along the floor with small movements of their feet.

Control-cost enhancement The first enhancement employed involved reducing the control cost within the fitness function. To encourage gaits that use more motion in novel ways, we applied a scalar multiplier to the control cost term, allowing for new behaviours at the cost of generating less efficient walkers (when the multiplier is below 1). This multiplier can be applied throughout each simulation episode or from a set timestep during each simulation episode. The latter was thought to be potentially beneficial as the gaits evolved typically have a "catch" phase in which walkers align themselves from the starting position into cyclic motions, which can be disrupted by a low control cost. The aim of this set of runs was to produce novel gaits with longer walks by reducing the control cost in the fitness function.

Balance enhancement The second enhancement employed was a fail condition involving the balance of the walker. In the original system, balance is described as the walker's torso's center's vertical (z-) component being outside the range of 1-2 simulation units. To improve walkers' episode durations (walking times) an additional constraint was introduced for the x- and y-dimensions, to terminate walkers with less upright postures. If the torso's center of mass moved outside a circle centered at the midpoint between the walker's two feet (each projected down to the ground plane) then it was considered a failure. The circle's radius is set as a multiple of the current distance between the midpoint and (either) foot. The radius-multiplier can be chosen to give a more or less tight constraint. The aim of this set of runs was to produce novel gaits with longer walks through stricter balance enforcement.

Combined enhancement The two enhancements were also combined, using the most successful parameters (control-cost multiplier and its delay, and balance-circle radius-multiplier) for each. The aim of this was to test if such a combination would achieve superior results.

Robustness to Action Noise

The MuJoCo 3D Humanoid-v1 environment contains a parameter for the standard deviation (*ac_noise_std*) of Gaussian noise to be added to the actions taken by walkers. To test the robustness of the evolved gaits, we evaluated evolved walkers with noise levels (*ac_noise_std*) from 0 (no noise) to 1, in order to observe the degradation of each major metric (average speed, distance travelled and episode time) until the walkers no longer achieved (lengths long enough to be typical of) stable gaits. The aim of this was to test whether or not the combined enhancement would result in more robust gaits, i.e. gaits with higher values in these metrics at higher levels of noise.

Results

Control-cost enhancement

Figure 2 shows the results for the default (d, no-enhancement) evolved behaviours and for walkers evolved with a control-cost multiplier of 0.25. 0.25 was chosen due to its promising results from initial testing of several multiplier values from 0.2 to 5. The multiplier was applied throughout (xp25) or from timestep 150 (s150, an estimate of the time at which the successful default gaits were reaching cyclic motion) or from timestep 500 (s500, the halfway point for a full-length episode). Median speed (averaged over the time of each evaluation) in the s500 runs was more than twice that in the default runs, with the former distribution significantly higher than the latter (Mann-Whitney $U=112$, $n_1=n_2=20$, $p<0.05$ one-tailed). Median distance traveled in the s500 runs was also more than twice that in the default runs, with the former distribution significantly higher than the latter (Mann-Whitney $U=121$, $n_1=n_2=20$, $p<0.05$ one-tailed). For time per episode, all medians were the maximum value (1000) and no significant difference was found.

Figure 5 includes high-performing gaits produced by the default and s500 runs, without action noise. The default runs' gait (top-left) shows a shuffling behaviour based around the knee joints. The s500 runs' gait (top-right) also shows a shuffling gait using the knee joints but, unlike the gaits produced by the default, the knees cross over, putting one leg in front of the other. This improved gait may be due to reduced importance of keeping energy expenditure low (at least per timestep rather than per unit distance) once a walker has reached a cyclic motion.

Balance enhancement

Figure 3 shows the results for the default (d) evolved behaviours and for walkers evolved with the balance enhance-

ment, with balance-circle radius-multipliers 1.00 (rp100), 0.75 (rp75), 0.5 (rp50) and 0.25 (rp25). Median speeds in the rp75 and rp50 runs were more than twice that in the default runs, with the former distributions significantly higher than the latter (rp75 Mann-Whitney $U=97$, $n_1=n_2=20$, $p<0.05$ one-tailed; rp50 Mann-Whitney $U=92$, $n_1=n_2=20$, $p<0.05$ one-tailed). Median distances traveled in the rp75 and rp50 runs were also more than twice that in the default runs, with the former distributions significantly higher than the latter (rp75 Mann-Whitney $U=91$, $n_1=n_2=20$, $p<0.05$ one-tailed; rp50 Mann-Whitney $U=99$, $n_1=n_2=20$, $p<0.05$ one-tailed). For time per episode, all medians were the maximum value (1000) except for the rp25 runs, which failed to produce a successful gait; the rp75 distribution was significantly higher than the default (Mann-Whitney $U=150$, $n_1=n_2=20$, $p<0.05$ one-tailed).

Figure 5 includes (bottom-left) a high-performing gait produced by the rp75 (0.75 radius-multiplier) runs, without action noise. The walker drags itself forward with one leg while pumping its arm for momentum, a behaviour previously unseen in the gaits produced by the default.

Combined enhancement

Figure 4 shows the results for the default (d) evolved behaviours, for walkers evolved with the most successful control-cost and balance-enhancement parameters (s500 and rp75) and for those evolved with the two enhancements combined (s500 combined with rp75). Median speed and median distance in the combined-enhancement were again more than twice those in the default runs, with the former distributions significantly higher than the latter (speed Mann-Whitney $U=79$, $n_1=n_2=20$, $p<0.05$ one-tailed, distance Mann-Whitney $U=74$, $n_1=n_2=20$, $p<0.05$ one-tailed). No significant increase in median speed or median distance was found between the s500 or rp75 runs and the combined-enhancement runs, although all four one-tailed p-values were below 0.15, suggesting that further runs might be useful to better investigate whether or not there is a significant difference here; however, it seems unlikely that any difference in median (between s500 and combo or between rp75 and combo) would be large even if significant. This may be because the two enhancements work in opposite directions, in that one is a reduction in control cost, allowing greater movement, while the other is a restriction on movement. For time per episode, the combined-enhancement runs matched the rp75 runs in having all gaits reach 1000 timesteps, although it is possible that increasing the 1000-timesteps limit could reveal a difference.

In the absence of noise, the highest-performing combined-enhancement gaits showed no noteworthy novelties, with all featuring either wide-legged shuffling with no leg crossover or single-leg dragging gaits. An example of the former can be seen in figure 5 (bottom-right).




Figure 2: Results for the default evolved behaviours and for walkers evolved with the 0.25 control-cost multiplier throughout, from 150 timesteps and from 500 timesteps: speed (left), distance (middle) and time (right) for each set of 20 runs.




Figure 3: Results for the default evolved behaviours and for walkers evolved with balance-circle radius-multipliers 1.00, 0.75, 0.5 and 0.25: speed (left), distance (middle) and time (right) for each set of 20 runs.




Figure 4: Results for the default evolved behaviours and for walkers evolved with the 0.25 control-cost multiplier from 500 timesteps, with the 0.75 balance-circle radius-multiplier and with a combination of the two: speed (left), distance (middle) and time (right) for each set of 20 runs.



Figure 5: Gaits without noise. Frame order is from left to right and top to bottom. In a high-performing gait from the default runs (top-left) the biped shuffles by alternating its knees in an unnatural motion. In a high-performing gait from the s500 runs (0.25 multiplier from 500 timesteps, top-right) the biped puts one leg in front of the other in succession, with a much wider range than the default's shuffling behaviours. In a high-performing gait from the rp75 runs (bottom-left) the biped pulls itself forward on one leg and pumps one arm for momentum, something previously unseen. In a high-performing gait from the combo runs (bottom-right) the biped shuffles similarly to the default gait, but with a wider spread of the legs. (The darkness at the end of the figures occurs as the humanoids walk out of the range of the white floor texture.) Videos of these gaits can be found at <https://github.com/KeeleBenJa/bipedal-methods>.




Figure 6: Results from scaling action noise from 0 to 1 for the default, 500 delay, 0.75 radial multiplier and combination of the two: speed (left), distance (middle) and time (right) averaged over each set of 20 runs.

Robustness to Action Noise

Figure 6 shows the degradation of average speed, distance travelled and time per episode with increasing levels of action noise. The combined-enhancement runs were much more robust to action noise than the default, particularly around noise=0.4. The rp75 and s500 runs produced intermediately robust gaits (i.e. more so than default but less so than for the combined-enhancement) at low levels of action noise but performed no better than default at medium to high levels of noise.

Figure 7 provides a closer look at the noise results at the 0.4 level, at which the combined-enhancement runs exhibited consistently superior results. In contrast to the without-noise results, the combined-enhancement results now show large, statistically significant improvements over the individual control-cost and balance enhancements. Median speed in the combined-enhancement runs was more than twice that in the default runs and the combined-enhancement distribution was significantly higher than the default runs (Mann-Whitney $U=79$, $n_1=n_2=20$, $p<0.05$ one-tailed), the s500 runs ($U=130$, $p<0.05$ one-tailed) and the rp75 runs ($U=99$, $p<0.05$ one-tailed). Median distance in the combined-enhancement runs was more than twice that in the default, s500 and rp75 runs, and the combined-enhancement distribution was significantly higher than the default runs (Mann-Whitney $U=53$, $n_1=n_2=20$, $p<0.05$ one-tailed), the s500 runs ($U=88$, $p<0.05$ one-tailed) and the rp75 runs ($U=70$, $p<0.05$ one-tailed). Median episode time in the combined-enhancement runs was also more than twice that in the default, s500 and rp75 runs, and the combined-enhancement distribution was significantly higher than the default runs (Mann-Whitney $U=67$, $n_1=n_2=20$, $p<0.05$ one-tailed), the s500 runs ($U=74$, $p<0.05$ one-tailed) and the rp75 runs ($U=62$, $p<0.05$ one-tailed).

Figure 8 shows the four previous high-performing gaits under noise level 0.4. The default (top-left), s500 (top-right, 0.25 control-cost multiplier from 500 timesteps) and combined-enhancement (bottom-right) runs produced similar gaits to figure 5 but with much wider motions, flailing limbs more. The rp75 (bottom-left, 0.75 radius-multiplier) gait arches the walker's back a little more but otherwise remains more stable, continuing to pump its arm.

Conclusions

Two fitness-function enhancements were tested to assess their affects on the speed, distance and duration of 3D Humanoid-v1 walks evolved using a replica of Salimans *et al.*'s evolution strategy (Salimans *et al.*, 2017). The first enhancement was to reduce control cost within the fitness function. When control cost was reduced to a quarter of the default, from the 500th timestep (the halfway point for a full-length episode), median speed and median distance both doubled. The most notable gait produced using this enhancement had a more pronounced stance and swing phase,

putting one leg in front of the other in a clearer fashion than any gait produced using the default fitness function.

The second enhancement was a balance enhancement, terminating walking when the torso's center of mass moved outside a circle centered at the midpoint between the walker's feet. When the circle's radius was 0.75 times current distance between the midpoint and (either) foot, median speed and median distance both doubled. In the most notable gait produced using this enhancement, the walker drags itself forward with one leg while pumping its arm for momentum, a behaviour previously unseen in the gaits produced by the default fitness function.

The two enhancements were also combined, using the most successful parameters for each (as above). Median speed and median distance were again double those in the default runs. However, there was little difference in these medians between the combined-enhancement and individual-enhancement runs and, in the absence of noise, the combined-enhancement gaits showed no noteworthy novelties.

To test the robustness of the evolved gaits, we evaluated evolved walkers with the addition of noise to their actions. The combined-enhancement gaits were much more robust to action noise than the default. The individual-enhancement gaits were intermediately robust at low levels of action noise but performed no better than default at medium to high levels of noise. In contrast to the without-noise results, the combined-enhancement gaits showed large, statistically significant improvements over those from the individual control-cost and balance enhancements, with median speed, distance and time two to five times those of the default and individual-enhancement gaits. This shows that the two enhancements synergise well to produce gaits that are robust to noise in their actions.

Future work will include further investigation into the reason for the large improvements exhibited by the combined-enhancement gaits in the presence of action noise, in contrast to the without-noise results. Removing or increasing the 1000-timesteps episode time limit would help to establish how much further the combined-enhancement can improve results, with and without noise. This will also help in establishing an upper bound for the optimal delay for the control-cost enhancement. OpenAI have ported the Humanoid walking task (and other tasks) from MuJoCo to the open-source Bullet Physics Engine and added more challenging tasks involving running toward a flag (<https://openai.com/blog/roboschool/> [https://github.com/openai/roboschool/](https://github.com/openai/roboschool)). Despite the emergence of improved locomotive gaits, the resulting controllers do not yet produce gaits as lifelike as that demonstrated in (Reil and Husbands, 2002), or achieve as complex behaviour as that demonstrated in (Heess *et al.*, 2017). We intend to evaluate our individual and combined enhancements on these more challenging tasks and to investigate the use of




Figure 7: Results with action noise=0.4; for the default evolved behaviours and for walkers evolved with the 0.25 control-cost multiplier from 500 timesteps, with the 0.75 balance-circle radius-multiplier and with a combination of the two: speed (left), distance (middle) and time (right) for each set of 20 runs.



Figure 8: Gaits with noise = 0.4, from the controllers shown (without noise) in figure 5. Frame order is from left to right and top to bottom. In the default-runs gait (top-left) the biped shuffles by alternating its knees in a more erratic way than before. In the s500-runs gait (top-right) the biped still puts one leg in front of the other, but much more loosely. In the rp75-runs gait (bottom-left) the biped uses one arm for momentum and is not affected too heavily by the noise, only bending its back more. In the combo-runs gait (bottom-right) the biped shuffles on its knees with a wide stance, making more flailing motions than previously. (The darkness at the end of the figures occurs as the humanoids walk out of the range of the white floor texture.) Videos of these gaits can be found at <https://github.com/KeeleBenJa/bipedal-methods>.

deeper networks in all of these tasks.

Acknowledgements

Our code is based on code from Uber AI Labs (Petroski Such et al., 2018; Conti et al., 2018, <https://github.com/uber-research/deep-neuroevolution>), which was in turn based on code from OpenAI (Salimans et al., 2017, <https://github.com/openai/evolution-strategies-starter>).

We thank these researchers and organisations for making their code available. Our code can be found at <https://github.com/KeeleBenJa/bipedal-methods>.

References

- Beyer, H.-G. and Arnold, D. V. (2001). Theory of evolution strategies – a tutorial. In *Theoretical aspects of evolutionary computing*, pages 109–133. Springer.
- Bramble, D. M. and Lieberman, D. E. (2004). Endurance running and the evolution of Homo. *Nature*, 432(7015):345.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). OpenAI Gym. *Preprint arXiv:1606.01540*.
- Conti, E., Madhavan, V., Petroski Such, F., Lehman, J., Stanley, K. O., and Clune, J. (2018). Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In *Advances in Neural Information Processing Systems 31 (NIPS 2018)*, pages 5027–5038. Curran Associates, Inc.
- Emanuel Todorov, E., Erez, T., and Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE.
- Fukuda, T., Komata, Y., and Arakawa, T. (1997). Stabilization control of biped locomotion robot based learning with GAs having self-adaptive mutation and recurrent neural networks. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 217–222. IEEE.
- Gökçe, B. and Akin, H. L. (2010). Parameter optimization of a signal-based omni-directional biped locomotion using evolutionary strategies. In *Robot Soccer World Cup*, pages 362–373. Springer.
- Guertin, P. A. (2013). Central pattern generator for locomotion: anatomical, physiological, and pathophysiological considerations. *Frontiers in neurology*, 3:183.
- Heess, N., Sriram, S., Lemmon, J., Merel, J., Wayne, G., Tassa, Y., Erez, T., Wang, Z., Eslami, A., Riedmiller, M., et al. (2017). Emergence of locomotion behaviours in rich environments. *Preprint arXiv:1707.02286*.
- Hewes, G. W. (1961). Food transport and the origin of hominid bipedalism. *American Anthropologist*, 63(4):687–710.
- Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks*, 21(4):642–653.
- Lehman, J., Chen, J., Clune, J., and Stanley, K. O. (2018). ES is more than just a traditional finite-difference approximator. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 450–457. ACM.
- Petroski Such, F., Madhavan, V., Conti, E., Lehman, J., Stanley, K. O., and Clune, J. (2018). Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *Preprint arXiv:1712.06567*.
- Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168.
- Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *Preprint arXiv:1703.03864*.
- Taga, G., Yamaguchi, Y., and Shimizu, H. (1991). Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological cybernetics*, 65(3):147–159.
- Tassa, Y., Erez, T., and Todorov, E. (2012). Synthesis and stabilization of complex behaviors through online trajectory optimization. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4906–4913. IEEE.
- Van der Noot, N., Ijspeert, A. J., and Ronsse, R. (2015). Biped gait controller for large speed variations, combining reflexes and a central pattern generator in a neuromuscular model. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6267–6274. IEEE.
- Vukobratović, M. and Borovac, B. (2004). Zero-moment point – thirty five years of its life. *International journal of humanoid robotics*, 1(01):157–173.
- Winter, D. A. (1991). *Biomechanics and motor control of human gait: normal, elderly and pathological*. University of Waterloo Press.