

A Simple 3D-Only Evolutionary Bipedal System with Albatross Morphology for Increased Performance

Ben Jackson, Alastair Channon
School of Computing and Mathematics
Keele University
ST5 5BG, UK

b.jackson@keele.ac.uk, a.d.channon@keele.ac.uk

Abstract—Bipedal walking is a difficult behaviour to encode into an evolutionary neural network, particularly in three-dimensional environments. Agents must be constantly maintaining balance alongside their primary objectives. Here we re-implement a simple evolutionary bipedal system, achieving high fitness and stepping gaits in 3D without the preliminary 2D bootstrapping process required by the original work. This high-performing system, with its deliberately simple neurocontroller, provides an excellent foundation for the community to use for the evolution or learning of more complex behaviours in bipeds. We also investigate the effects of modified morphology with the system, significantly improving agent fitness by evolving networks alongside morphologies resembling a baby albatross. The agents with albatross morphologies travel up to three times further than default agents. We then test incrementally evolving agent morphology via the simultaneous evolution of a separate morphological genotype. We initialised this genotype either alongside a high-performing controller or from a completely random point in both fitness landscapes. Agents evolved from this random initialisation travel up to four times further than default agents. One randomly initialised incremental morphology also achieves gaits with significantly higher upper body and swing knee controller input weights than the default.

1. Background

Bipedal locomotion has been the focus of many studies [22]. Its emergence has been linked with endurance running [2] and tool use [11]. The main challenges for bipedal agents are balancing and generating cyclic motion across two limbs. In robotics, bipedal walking is often achieved through zero-moment point computation [21], for evaluating stability [8].

Initially, bipedal systems utilised Central Pattern Generators (CPGs) as a control network for neuroevolution. CPGs are recurrent neural networks that produce rhythmic output from a linear input set [10]. Modelled after pattern generators found in human and animal spines, they offer a biologically inspired solution; demonstrated in [19], [12] and [20]. Reil *et al.* used a genetic algorithm to evolve

CPG network weights, producing a stable bipedal walking agent [15]. Measuring fitness as distance from the point of origin, they achieved a 10% stability success rate. This increased to 80% through the introduction of an oscillation bonus in the fitness function. Reil *et al.* added sensors to achieve turning behaviours. Turning was also achieved in [9], using evolution strategies [1] to produce efficient CPG-based walking behaviour, both in simulated and real-world bipedal agents.

In more recent works, bipedal agents have used simpler control schemes than the CPG. Solomon *et al.* used a genetic algorithm with a simple network to evolve bipedal gaits for rough terrain walking [18]. Their agents contained multiple perceptron-like ‘Linear Reactive Control’ networks as controllers, seen in Figure 1. One of these networks separately controlled each actuator. Solomon *et al.* evolved these controllers incrementally, seeding each stage of evolution with the fittest agents from the previous stage. In the initial stage, only angular inputs for the joint being actuated were fed to each network, but later configurations fed every joint input to every network. The end result was stable 3D walking on flat terrain with a 12% fall rate and a rough-terrain fall rate of 17%. Despite the simplicity of this system, early controllers had to be evolved on a 2D model, before being seeded into a 3D model. Here we will re-implement the system seen in [18] directly into 3D, without prior 2D processing. We hope this simplified system could provide a common starting point for future evolutionary bipedal work in the scientific community.

Morphology also plays a role in agent locomotion. McGeer designed a walker which travelled down a slope using gravity as its sole source of energy, moving its legs in a natural swinging motion [14]. Dubbing this a “passive dynamic walker”, it highlighted the potential of replacing sections of a controller with morphology. Collins *et al.* added passive swinging arms and knees to robotic walkers; demonstrated in [5] and [6]. Through this methodology, Collins *et al.* evolved stable gaits for both in-phase and anti-phase swinging. They also tested binding human subjects’ arms during walking, which demonstrated higher upper body momentum when unbound. In non-

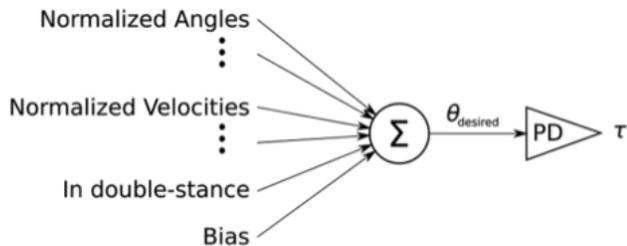


Figure 1. A linear reactive controller, from [18]. Solomon *et al.*'s walkers contained a separate controller for each actuator, producing a desired angle value from a linear vector of inputs. The desired angle value was then fed into a Proportional Derivative controller to produce actuator torque value.

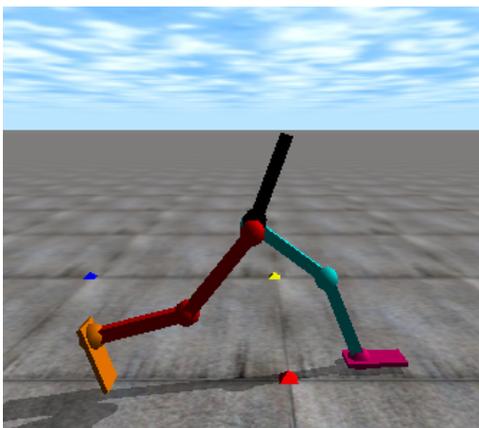


Figure 2. Our re-implemented walker in 3D, based on [18].

bipedal agents, Calisti *et al.* investigated the effects of mid-lifetime morphological changes on locomotion [3]. They employed a soft, underwater, octopus-like robot, as an underwater environment could accommodate walking, running and swimming gaits; this allowed for easier transitions. They produced multiple distinct gaits which could transition into each other. This included walking, running, and sculling, thanks to the robot's soft appendages.

The baby albatross has a set of fluffy feathers on its underbelly that soften impacts and allow it to roll upright to regain balance. Feathers are theorised to have emerged from scales in order to repel water and produce an airtight barrier [7]. Their softness was also proven to assist with adjusting flight trajectory, both in simulated [16] and real world [4] agents. We will apply albatross-based morphologies to our re-implementation and demonstrate the increase in fitness they provide. We will also investigate incremental morphological evolution: seeding new agent populations with the controllers of previous winners, in order to simultaneously evolve morphologies capable of more human-like gaits. In order to assess these gaits, we will analyse the genotypes of our successful agents, observing which has the highest input weights for each controller.

2. Methods

2.1. Re-implementation

Solomon *et al.* [18] used a linear reactive network to produce a torque value for each actuator, seen in Figure 1. Each agent contained 2 thigh segments, 2 shank segments, 2 feet segments and an upper body segment. Actuators at the hip, knee and ankles joined these segments together. The 3D model featured an extra lateral hip actuator in the centre of the upper body and passive lateral ankle joints. Solomon *et al.* fed a set of inputs to each controller: joint angles, joint angle differentials (also referred to as velocities in Figure 1), a flag for both feet touching the floor, and a bias. These inputs were then normalised and multiplied by a set of weights; they evolved these weights as the agent's genotype. After this, they summed the resulting values to form the desired output angle values. The desired angle values were then fed through a Proportional Derivative (PD) controller [13], modifying the angle values in response to the size and derivative of their previous error; this produced final torque values for the actuators. Solomon *et al.* tested several control variations: Local Proportion, where only angles and derivatives for the actuator's joint were fed to each network; Fully Connected, where every joint angle and derivative was fed to each network; and Reduced Sparsely Interconnected, where the networks were pruned down from a fully connected state until fitness sufficiently degraded. They evolved the first two network configurations in a 2D physics engine; successful controllers from their populations were then seeded into the 3D engine and pruned down. Solomon *et al.* measured fitness as distance travelled forwards. Termination occurred when an agent either fell over, segments above the thigh contacted the ground, or a lifetime torque allowance across all actuators was exceeded. In this way, efficient gaits were promoted which were more likely to resemble natural gaits; energy efficiency is important for survival in most species. In each generation, the algorithm selected the top 20% of agents by fitness to be parents. Their offspring were then produced through mutation, single-point pointwise crossover or both. Each network had its own mutation rate which was evolved as part of the genotype, initialised at a value of 0.1. Elitism retained the best-performing agent of each generation. This allowed for better navigation of the rugged bipedal fitness landscape. The top 20% of agents in the last 10% of generations were compared to select a winner. This method recovered lost behaviours from penultimate generations that were potentially more successful. We chose Solomon *et al.*'s work for re-implementation as the rough terrain results demonstrated robust behaviour, using a deliberately simple controller for bipedal locomotion. The aim of our work was to develop this into a better foundation for the future evolution of more complex bipedal behaviours.

In order to evolve the system in 3D from scratch with the simple linear reactive control scheme, we modified the existing system seen in [18]. We created our re-

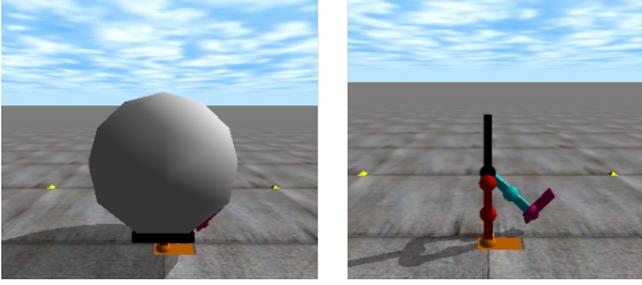


Figure 3. (left) The square base morphology with the square-based sphere visible. (right) The same walker without, demonstrating the shorter legs.

implementation with the Open Dynamics Engine [17]; it featured bipedal agents configured with identical body parameters to the original work, seen in Figure 2. The genetic algorithm used was also identical, with a population size of 150 for 500 generations. The agents contained six linear reactive controllers, with inputs configured identically to the Partially Connected controllers described above. Unlike the original networks, our linear networks did not contain PD controllers. Instead, the desired angle values were multiplied by a speed constant and then set as the velocity parameter for each actuator. There were also no lateral joints in our 3D model, in order to reduce the size of the fitness landscape. Agents were initialised standing on one leg with the other swinging up, away from the ground. By starting in a swinging motion, they could transition into a cyclic gait with less difficulty. We replaced torque limits with a 60 second lifetime; a flight phase, both feet being off the ground, was also permitted. This allowed the agents to exhibit more motion rather than prioritising efficiency. We also applied 20% elitism. By applying these modifications to the system, we were able to evolve successful bipedal agents entirely in 3D.

2.2. Albatross Morphology

2.2.1. Square Base. The first albatross morphology tested is referred to as the square base morphology. The morphology featured a large sphere segment attached to the upper body, with a lower cuboid segment to give the sphere a square base. This would resemble the lower body of a baby albatross. The legs were also scaled to half length to allow the sphere to contact the ground. Shown in Figure 3, we aimed for this morphology to allow bipedal agents to roll back onto the base using the sphere; after this the square base would allow them to realign themselves, keeping the legs parallel with direction of travel. To achieve this, the sphere was exempt from the system’s fall condition when it contacted the ground. Whilst this technically prevented the bipedal agent from falling, it did not guarantee higher fitness or longevity.

2.2.2. Big-Tucked. The second morphology tested is referred to as the big-tucked morphology. Shown in Figure 4,

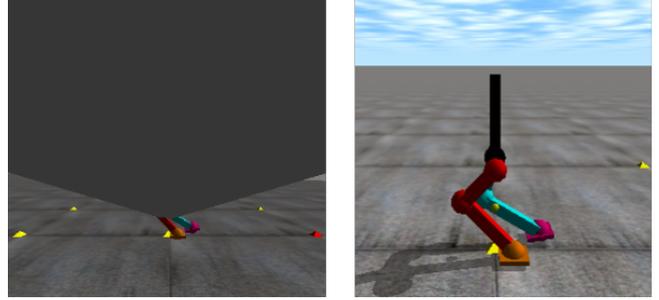


Figure 4. (left) The big-tucked morphology with the big sphere visible. (right) The same walker without, highlighting the inverse knees, shortened feet and white counterweight between the legs.

it featured a vastly larger sphere attachment, without a square base. A small spherical counterweight was also added inside the larger sphere. We hoped these would enable agents to roll back upright, much like the square base morphology. As well as this, we also modified the knees and feet. The baby albatross exhibits legs that appear to bend inwards, to fold away when they sit; in fact this is an ankle joint, with an elongated foot. We inverted our agents’ knee joint limits to resemble this. Aiming for this to allow agents to similarly tuck their legs away into the larger sphere, we also hoped it would allow them to lift themselves off the ground, akin to a scissor lift. The feet were shortened to prevent the agents from using the tucked legs to drag themselves along instead of stepping. This morphology was also not able to fall due to the larger sphere’s contact exemption.

2.2.3. Incremental Morphological Evolution. We also tested incrementally evolving the two albatross morphologies. This involved seeding new agent populations with the genotypes of high-performing agents from previous runs. To better achieve this, we created a separate genotype for the agents’ morphology. This featured two genes: a single multiplier value, applied to both the upper and lower leg lengths, and a mutation rate. The multiplier gene value was limited to being between the albatross leg length (multiplier value of 1) and the default length (multiplier value of 2). The body genotypes were simultaneously mutated and crossed over with the same methodology as the controller genotypes. We aimed to take the shorter-legged albatross morphologies and evolve them back to default proportions, whilst retaining their increased performance. We tested evolving agents with seeded (referred to in later sections as preloaded) genotypes from high-performing albatross agents, and also agents with both genotypes randomly initialised.

3. Results

3.1. Re-implementation

We evolved each tested morphology in the re-implemented system as a set of twenty runs, measuring

fitness as distance travelled forwards from the origin. Agents failed when any part of the body above the knee but besides the sphere attachments contacted the ground (a fall), or they exceeded a 60 second lifetime. Compared to the original work, our default morphology agents did not perform as well. They were not able to walk on rough terrain, or travel similar distances. Despite this, as seen in Figure 9, they were still able to achieve walking gaits. They also still travelled for some distance before falling over. Agent legspan was around 1.2, and Figure 5 shows distances travelled as high as 30. As well as this, all of the gaits in this work were evolved in 3D from scratch with no prior processing. We suspect the main reason for the decreased performance was the lack of 3D actuators (passive ankle tilt and lateral hip actuation) seen in the original work’s 3D model. More lateral control would allow for increased balance, to avoid falling sideways during a stepping motion.

3.2. Square Base and Big-Tucked

Figure 5 shows the fitness of the twenty winning default, square base and big-tucked morphologies. The square base had a much higher median than the default, but a larger range. The big-tucked morphology had a slightly larger median than the square base, and a shorter range. Both the square base (Mann-Whitney $U=10$, $n_1=n_2=20$, $p=1.008e-09$ one-tailed, without correction for multiple comparisons) and the big-tucked (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons) were significantly fitter than the default. Figure 6 shows the average fitness curves for twenty of each morphology. The square base and big-tucked morphologies evolved at a slower pace but achieved final fitness values around three times higher than the default.

3.3. Incremental Square Base and Big-Tucked

Figure 7 shows the fitness of twenty winning agents for the default and each incremental morphology type. These included square base preloaded, square base randomly initialised, big-tucked preloaded and big-tucked randomly initialised morphologies. The square base preloaded (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons), square base randomly initialised (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons), big-tucked preloaded (Mann-Whitney $U=36$, $n_1=n_2=20$, $p=6.835e-07$ one-tailed, without correction for multiple comparisons), and big-tucked randomly initialised (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons) were all significantly fitter than the default. The square base randomly initialised had a higher median, but also a larger range than the preloaded version. The square base randomly initialised morphology was also significantly fitter than the preloaded version (Mann-Whitney $U=14$, $n_1=n_2=20$, $p=3.685e-09$ one-tailed, without correction for multiple comparisons). The big-tucked morphologies had similar range but drastically

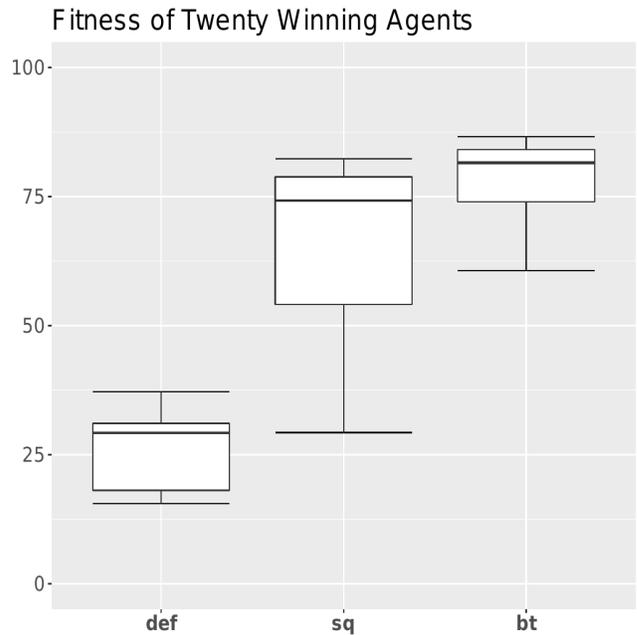


Figure 5. The fitness of twenty winning agents for each set of morphologies. Def represents the default morphology, sq represents the square base morphology and bt represents the big-tucked morphology.

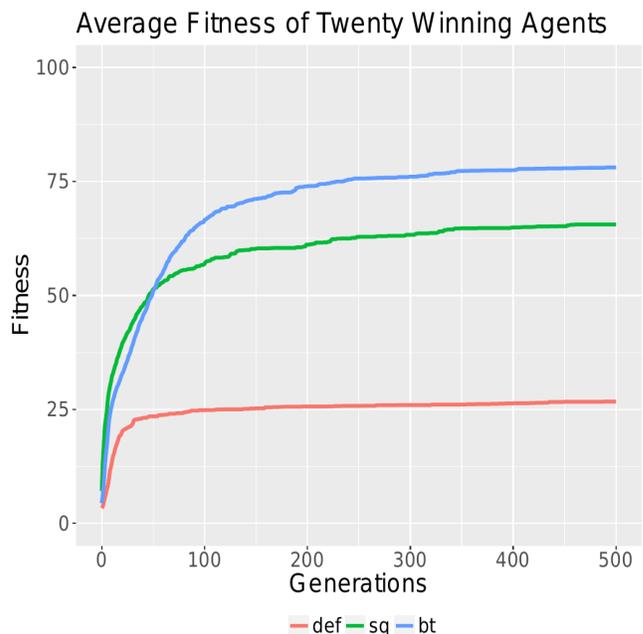


Figure 6. The average fitness curves of twenty winning agents for each set of morphologies. Def represents the default morphology, sq represents the square base morphology and bt represents the big-tucked morphology.

different medians; the randomly initialised morphology was significantly fitter (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons). The big-tucked randomly initialised median was more than three times that of the preloaded version, and more than four times the median of the default version. Figure 8 shows the average fitness curves for these morphologies. All the morphologies evolved at a similar pace, with the big-tucked evolving slightly faster than the square bases. All finished above the default, with the randomly initialised versions outperforming the preloaded versions on both morphologies. The difference in performance between the two big-tucked morphologies was much greater than that of the square bases. Finally, we recorded gait snapshots from high-performing agents in the two fittest morphology groups, both randomly initialised morphologies, and the default. We took each gait snapshot over the same time period, about five seconds, at the same point during each agent’s lifetime, in order to capture the default’s fall. Figure 9 shows a gait from the default morphology. It exhibits a stepping behaviour, but it is slow and unstable, eventually falling. Figure 10 shows a gait from the square base randomly initialised morphology. It travels forwards quickly in a can-can like motion. Its body has not moved away from that of the standard square base morphology. Figure 11 shows a gait from the big-tucked randomly initialised morphology. It travels forwards at a much higher speed than the default. It evolved much longer legs than the standard big-tucked morphology; these were closer to the default in length.

3.4. Genotype Analysis

To assess which morphologies were exhibiting the most human-like gaits, we analysed their controller genotypes. As described previously, the agents’ control networks evolved by our system consisted of six linear controllers, for straightforward reproduction. The agents’ controller genotype was represented by a set of network weights for these controllers. The controllers produced desired angle values for the upper body, interleg, stance knee, swing knee, stance ankle and swing ankle angles, which were then applied as velocity parameters to the actuators. The stance leg was the leg balanced upon whilst the swing leg swung forwards during a gait. These roles then alternated as agents took steps. The configuration we used only fed angle/derivative inputs to each controller that matched the joint it was actuating. Taking their absolute values, we calculated the mean of each input weight from the twenty winners for each morphology. This allowed us to observe which morphologies’ input weights had the highest mean values, and therefore how much each morphology utilised each controller.

Figure 12 shows the mean input weight values for the default, square base randomly initialised and big-tucked randomly initialised morphologies. The square base and big-tucked randomly initialised morphologies evolved agents with higher upper body input weights than the

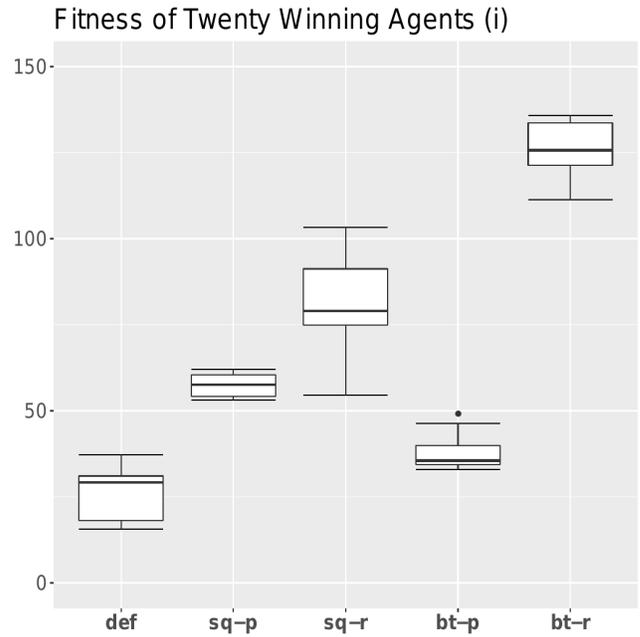


Figure 7. The fitness of twenty winners for each of the incremental morphologies and the default. Def represents the default, sq represents the square base morphology and bt represents the big-tucked morphology. A suffix of p represents a preloaded morphology. A suffix of r represents a randomly initialised morphology.

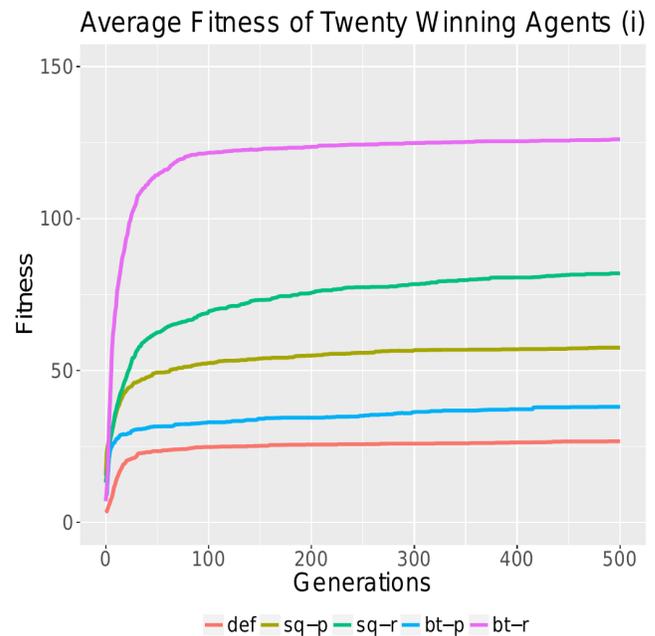


Figure 8. The average fitness curves of twenty winners for each of the incremental morphologies and the default. Def represents the default, sq represents the square base morphology and bt represents the big-tucked morphology. A suffix of p represents a preloaded morphology. A suffix of r represents a randomly initialised morphology.

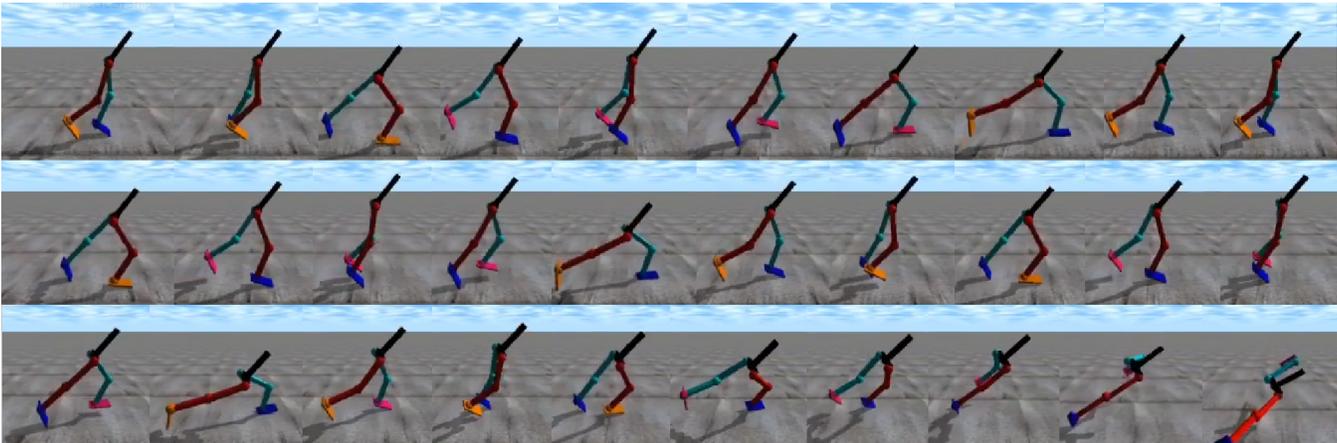


Figure 9. (left to right, top to bottom) A high performing gait from the default morphology.

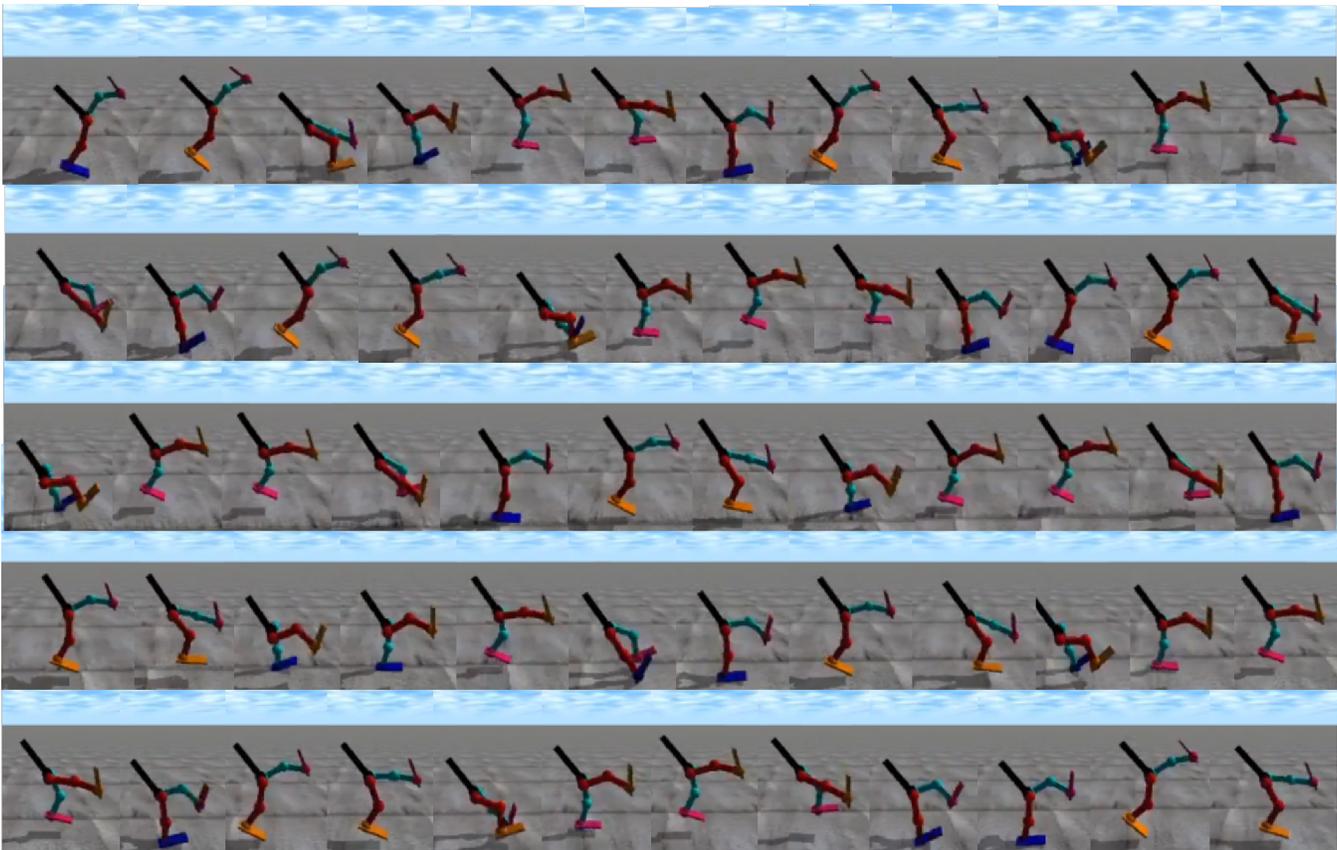


Figure 10. (left to right, top to bottom) A high-performing gait from the square base randomly initialised morphology. The square based sphere attachment is not drawn.

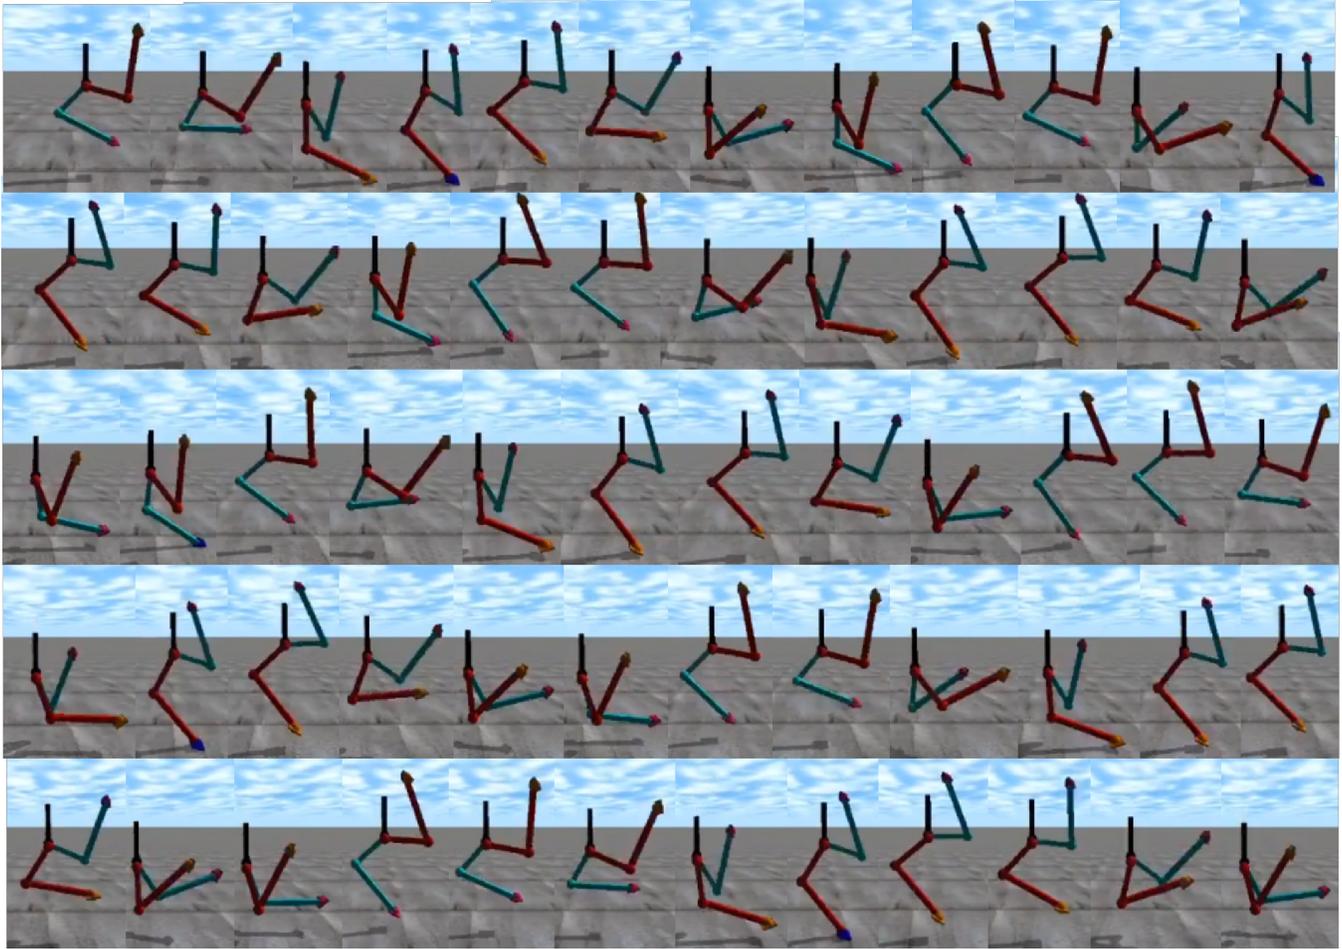


Figure 11. (left to right, top to bottom) A high performing gait from the big-tucked randomly initialised morphology. The large sphere attachment and counterweight are not drawn.

default. This aligned with the square base and big-tucked randomly initialised gaits lasting longer than the default gait. We also compared the morphologies' input weights directly without averaging, which demonstrated the square base randomly initialised upper body angle (Mann-Whitney $U=1$, $n_1=n_2=20$, $p=1.451e-11$ one-tailed, without correction for multiple comparisons) and upper body derivative (Mann-Whitney $U=0$, $n_1=n_2=20$, $p=7.254e-12$ one-tailed, without correction for multiple comparisons) were significantly higher than the default. The big-tucked randomly initialised upper body angle (Mann-Whitney $U=1$, $n_1=n_2=20$, $p=1.451e-11$ one-tailed, without correction for multiple comparisons) and upper body derivative (Mann-Whitney $U=8$, $n_1=n_2=20$, $p=4.86e-10$ one-tailed, without correction for multiple comparisons) were also significantly higher than the default. Figure 12 shows the big-tucked randomly initialised morphology had higher swing knee input weights than the default, with significantly greater values for swing knee angle (Mann-Whitney $U=48$, $n_1=n_2=20$, $p=5.68e-06$ one-tailed, without correction for multiple comparisons) and swing knee derivative (Mann-Whitney $U=48$, $n_1=n_2=20$,

$p=5.68e-06$ one-tailed, without correction for multiple comparisons). This aligned with the randomly initialised big tucked gait swinging its legs faster than the default.

4. Conclusions

We produced a functional re-implementation of [18] directly into 3D without bootstrapping from 2D first; featuring a deliberately simple linear control scheme in order to provide a foundation for future bipedal work both from ourselves and in the community. The re-implementation achieves high fitness and stepping gaits. We introduced two novel morphology styles based on a baby albatross to the re-implementation; one with a square base and the other with tucked knees. These morphologies were up to three times fitter than the default, despite not technically being able to fall. We then incrementally evolved these morphologies, creating a second morphological genotype and seeding both with previously high-performing agents. Randomly initialising these morphology and controller genotypes instead achieved fitness up to four times higher.

Default Genotype			Square Base Genotype			Big Tucked Genotype		
	angle	derivative		angle	derivative		angle	derivative
upper body	0.03898	0.02870	upper body	0.08215	0.07830	upper body	0.08375	0.07212
interleg	0.03702	0.03346	interleg	0.03332	0.03472	interleg	0.03602	0.03871
stance knee	0.02128	0.02921	stance knee	0.03051	0.04352	stance knee	0.03818	0.01666
swing knee	0.03908	0.03335	swing knee	0.01694	0.02955	swing knee	0.05347	0.05171
stance ankle	0.02491	0.02218	stance ankle	0.02696	0.03952	stance ankle	0.03068	0.03504
swing ankle	0.03216	0.03402	swing ankle	0.03133	0.02963	swing ankle	0.02830	0.02282

Figure 12. We list the agents’ six controllers, upper body to swing ankle, on the left of the tables. (left to right) The default, square-base randomly initialised and big-tucked randomly initialised morphologies’ mean input weights.

Finally, the big-tucked randomly initialised morphology evolved an extremely fast gait with significantly higher input weights for its upper body and swing knee controllers.

Future work will include incrementally evolving more complex bipedal behaviours, such as jumping or multi-directional running, from this simpler motion. We will also execute the system with the lateral joints added; this might help determine why the randomly initialised morphologies outperformed the preloaded ones.

References

- [1] Beyer, H-G., and Dirk V. Arnold. “Theory of evolution strategies—A tutorial.” *Theoretical aspects of evolutionary computing*. Springer, Berlin, Heidelberg, 2001. 109-133.
- [2] Bramble, Dennis M., and Daniel E. Lieberman. “Endurance running and the evolution of Homo.” *nature* 432.7015 (2004): 345-352.
- [3] Calisti, Marcello, et al. “An octopus-bioinspired solution to movement and manipulation for soft robots.” *Bioinspiration & biomimetics* 6.3 (2011): 036002.
- [4] Chang, Eric, et al. “Soft biohybrid morphing wings with feathers underactuated by wrist and finger motion.” *Science Robotics* 5.38 (2020).
- [5] Collins, Steven H., Martijn Wisse, and Andy Ruina. “A three-dimensional passive-dynamic walking robot with two legs and knees.” *The International Journal of Robotics Research* 20.7 (2001): 607-615.
- [6] Collins, Steve, et al. “Efficient bipedal robots based on passive-dynamic walkers.” *Science* 307.5712 (2005): 1082-1085.
- [7] Dyck, J. A. N. “The evolution of feathers.” *Zoologica Scripta* 14.2 (1985): 137-154.
- [8] Fukuda, Toshio, Youichirou Komata, and Takemasa Arakawa. “Stabilization control of biped locomotion robot based learning with GAs having self-adaptive mutation and recurrent neural networks.” *Proceedings of International Conference on Robotics and Automation*. Vol. 1. IEEE, 1997.
- [9] Gökçe, Barış, and H. Levent Akin. “Parameter optimization of a signal-based omni-directional biped locomotion using evolutionary strategies.” *Robot Soccer World Cup*. Springer, Berlin, Heidelberg, 2010.
- [10] Guertin, Pierre A. “Central pattern generator for locomotion: anatomical, physiological, and pathophysiological considerations.” *Frontiers in neurology* 3 (2013): 183.
- [11] Hewes, Gordon W. “Food transport and the origin of hominid bipedalism.” *American Anthropologist* (1961): 687-710.
- [12] Ijspeert, Auke Jan. “Central pattern generators for locomotion control in animals and robots: a review.” *Neural networks* 21.4 (2008): 642-653.
- [13] Johnson, Michael A., and Mohammad H. Moradi. *PID control*. Springer-Verlag London Limited, 2005.
- [14] McGeer, Tad. “Passive dynamic walking.” *I. J. Robotic Res.* 9.2 (1990): 62-82.
- [15] 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.
- [16] Shim, YoonSik, and Phil Husbands. “Feathered flyer: integrating morphological computation and sensory reflexes into a physically simulated flapping-wing robot for robust flight manoeuvre.” *European Conference on Artificial Life*. Springer, Berlin, Heidelberg, 2007.
- [17] Smith, Russell. “Open dynamics engine.” (2005): 3-2.
- [18] Solomon, Joseph H., Mark A. Locascio, and Mitra JZ Hartmann. “Linear reactive control for efficient 2D and 3D bipedal walking over rough terrain.” *Adaptive Behavior* 21.1 (2013): 29-46.
- [19] Taga, Gentaro, Yoko Yamaguchi, and Hiroshi Shimizu. “Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment.” *Biological cybernetics* 65.3 (1991): 147-159.
- [20] Van der Noot, Nicolas, Auke J. Ijspeert, and Renaud Ronsse. “Biped gait controller for large speed variations, combining reflexes and a central pattern generator in a neuromuscular model.” *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015.
- [21] Vukobratović, Miomir, and Branislav Borovac. “Zero-moment point—thirty five years of its life.” *International journal of humanoid robotics* 1.01 (2004): 157-173.
- [22] Winter, David A. *Biomechanics and motor control of human gait: normal, elderly and pathological*. 1991.