Evolving Sims's Creatures for Bipedal Gait

Amin Azarbadegan, Frank Broz and Chrystopher L. Nehaniv
School of Computer Science, University of Hertfordshire, UK
Email: {s.azarbadegan, f.broz, c.l.nehaniv}@herts.ac.uk

Abstract

In this paper we describe the design of an approach to evolve Sims's creatures with morphology and behaviour similar to biped animals. Our hypothesis is that biases in morphology that encourage limb specialisation, combined with rewards for successful locomotion and carrying at the same time and realistic, physics-based penalties for falling, would lead to creatures capable of bipedal locomotion. We present experimental results demonstrating successful evolution of biped morphology and stepping gait.

1. Introduction

Keeping dynamic balance and bipedalism, especially as seen in humans as an upright mode of locomotion, is a goal of robotics as a step towards making better humanoids. In robotics and control theory, walking gaits are typically generated without concern for whether they are natural and human-like. Human walking is adaptive and dynamic, while current non-tethered bipedal robots walk statically with predefined movements (for example, see [1]), yet these walking controllers are complicated and hard to design.

On the other hand, bipedalism has been relatively unexplored in artificial life research, while its natural evolution in hominids is a controversial question amongst the biologists and physical anthropologists. For example Vaughan et al evolved a dynamic walking controller which can adapt to external and internal noise [2]; but this system uses fixed morphology and neural network and focuses merely on the behaviour. Thus, we intended to utilise currently available AL techniques to synthesise an open evolutionary system where one can test some of the hypotheses about the evolution of bipedalism.

The humanistic form of bipedalism is not particularly beneficial in terms of efficiency and locomotory abilities. There is evidence that this form had emerged a few millions years before the hominids in an ape species, *Oreopithecus*, in insulated islands; but the species became extinct after the insulation finished, which may imply that this gait is not a winning factor in an open competition. There are hypotheses about thermoregulation benefits of humanistic bipedal gait in the open, unforested deserts but doubt has been cast on these after new discoveries suggesting the hominid origins might have been in the forests [3, pages 269-270].

However, evolutionary theorists agree that human bipedalism is favoured by natural selection because it freed hands for employing tools and/or carrying food for long distances, and also made humans fit for endurance running [4], [5].

2. Background

2.1. Sims's creatures

Karl Sims innovated a new way of evolving virtual artificial agents which generates their morphology and behaviour simultaneously in a simulated physical environment [6]. The genomes are encoded based on graph-rewriting (similar to L-systems [7]) allowing for the modular reuse of components, segmentation, and symmetry in the body plans, but also similarly in controlling computational elements within and connecting the different blocks making up the body allowing the creatures to sense and act. Sims later employed co-evolution through interactive competitions in one-on-one contests in which creatures compete for the control of a cube [8].

The Sims's creatures are composed of rigid blocks. Their movements are generated by the joints between the blocks and controlled by an internal neural network that receives input from sensors. A creature can be wholly and identically produced by its genetic code. Applying a genetic algorithm, after a few generations evolution of fitter creatures can be expected.

Sims's work was one of the first open evolutionary systems that generated virtual 3D agents in a simulated physical environment and initiated a series of similar research and simulations, of which an extensive and relatively recent review can be found in Miconi's thesis [9, page 34 ff.].

However, Miconi claims that until his work in the next decade no one has replicated Sims's efforts "to achieve results comparable to him in efficiency and complexity" [9]. Furthermore, Miconi used Newtonian physics, simulated by Open Dynamics Engine (ODE), in contrast to simplified laws of many other simulation environments, and used standard McCulloch-Pitts neurons instead of complex and functional ones as in the Sims model. He managed to breed creatures for locomotion and grabbing abilities (the Sims's original experiments) and went further by simulation of fighting contests.

2.2. Bipedalism in the Sims's creatures

The evolution of bipedal creatures is not observed in the Sims's experiments nor in their replication in Miconi's work, as the selection pressure was not towards any particular gait; rather, it was merely focused on efficient movement. Bipedalism has not been ubiquitous in the natural evolution of animals either; particularly as seen in humans with an upright mode of locomotion: bipedalism can be seen in dinosaurs, birds and kangaroos, but they all have their upper part of body balanced with a tail, which is only vestigial and internal in humans.

We can test hypotheses about the evolution of bipedalism in a physics-simulation artificial model, and see whether biped creatures emerge in a model that breeds them for their ability of successful locomotion and carrying at the same time.

2.2.1. Falling damage. Biomechanists describe human walking as controlled falling [10]. This suggests a derivative sophistication in the dynamic balance control - a behaviour evolved much earlier than walking itself, as in the presence of gravity loss of balance might cause serious harms to the creature.

But there is no deterrent for falling - no matter how severe the impact is - in Sims's original model or its replications. Contrarily, some creatures use this unrealistic simplification to generate locomotion strategies with coarse leaping gaits, whose unabsorbed shocks would damage creatures in the real world. Sims himself mentions the problem of an "inelegant strategy of falling" in his original paper and in the box-grabbing contest he had to put some limitations to discourage "strategies like this that utilize only potential energy" [8].

As he puts it "it is important that the physical simulation be reasonably accurate when optimizing for creatures that can move within it. Any bugs that allow energy leaks from non-conservation, or even round-off errors, will inevitably be discovered and exploited by the evolving creatures" [6]. For the emergence of dynamic balance behaviour it seems inevitably reasonable to consider in detail the possible degrees of damage from falling onto the ground.

2.3. Genealogy

Genealogical techniques, i.e. tracing parents and ancestors of the final resulted individuals, are relatively unused in the field of genetic algorithms and the analytical and statistical studies over the generation averages have been preferred. However, in some cases like ours the difficulty of the task makes the evolution of successful creatures a rare event, and thus the statistical analysis becomes meaningless while genealogical techniques still can provide a useful insight towards quantities that are hard to measure, e.g. the variability of population in high dimensional search spaces.

Genealogy has been mainly used in analytical and theoretical studies about the genetic algorithms (for example, see [11] and [12]) and rarely used in practice because of its high computational cost (see [13] and [14]). However, it can be limited to a few number of generations to reduce the cost, for example *fitness transmission*, a genealogical measure that Miconi introduced in [15], traces back ancestors to two levels.

The genealogy tracing process recursively generates a binary tree with leaves as ancestors. At the first ancestral level, i.e. the direct parents, there are two leaves, and in the nth level 2^n leaves are known as ancestors, but they could be repetitive at each level, because a single creature in a generation can contribute several times in making the next generation. The number of repeats of a leaf in a generation divided by total of number of leaves gives the *contribution* of that ancestor in the genome of the studied individual. This contribution measure does not necessarily represent the real contribution since the reproduction methods are not equally neutral towards genomes of two parents, but can be used as an estimation.

3. The model

The model used in these experiments is an extension of the model used by Miconi. The simulation environment is adapted from Miconi's simulation code, an open-source implementation based on Sims's research.

The main class is *Animat*, which has a *genome*, an array of *genes*. Each gene corresponds to a *limb* and has data about its dimensions, parent, number of recursions, place of the connection to the parent and whether it is reflected on the other side of the parent or not. The gene also has the neural information of the limb in an array of neurons. Each gene can be initialised completely randomly within predefined limits.

We added limb specification to the model so there can be *HEAD*, *TORSO* or simply *LIMB*. The number and average dimensions of the initialised limbs and their recursions and reflections correspond to their specification in a hypothetically ideal phenotype, e.g., a creature has only one head.

To evaluate fitness of each randomly/evolutionarily generated creature, we ran a limited-time simulation. We hypothesise that the evaluation should encourage bipedal locomotion and deter damages.

In Miconi's model, at each time step during simulations the internal McCulloch-Pitts neural network is updated two times. We later increased the number of updates to six. Some of the neurons are defined as sensors and fire in response to the environmental inputs, some of them are actuators and move the joints according to their input and the others are in the *hidden layer* and process data.

The acceleration is assumed to be the only relevant stimulus, so we defined accelerometer sensors as part of neural network of the creatures and the only sensor type available in the model.

3.1. Fitness function

The most investigative part of this work has been the definition and fine tuning of the fitness function. The sought-after characteristics were:

- **Quantification**: bipedalism and dynamic balance are not easily quantifiable concepts. The function should include incentives for the upright morphology and biped locomotion as well as deterrents for falling onto the ground.
- **Distinction**: on one hand, a picky and strict fitness function would result in indistinguishably bad scores for every creature and could not give a gradient to elevate in the fitness landscape, since the creatures at the beginning are unlikely to be even close to the desired morphology, let alone to the right behaviour. On the other hand, a lenient function could end up producing arbitrary configurations.
- **Adequateness**: as no function could possibly guarantee the desired morphology, it might trap the GA in a local optimum far from a biped creature. Therefore, like many other GA applications, the initial fine tuning of parameters is crucial for the GA convergence and reasonable performance.

3.1.1. Damage. During the simulation, when a collision is detected, the impact is considered as the change of direction of the speed vector, scaled by the mass of the body. Formally, if at time step t a limb λ with mass m and linear speed V_t collides with the ground, we define

$$I_{\lambda,t} = -mV_t \cdot V_{t-1} \tag{1}$$

$$D_{t,\lambda}(\theta) = \begin{cases} I_{\lambda,t} & I_{\lambda,t} > \theta \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

in which $D_{t,\lambda}(\theta)$ is linear damage on the limb λ at time step t where *damage threshold* is θ, a predefined positive number. A similar damage $\Delta_{t,\lambda}(\theta)$ is calculated for the change of angular speed ν_t, in which moment of inertia tensor replaces mass. The total damage on the limb λ over n time steps is the sum of all damages, which is by construction guaranteed to be non-negative:

$$D_\lambda(\theta, n) = \sum_{i=0}^{n} D_{t,\lambda}(\theta) + \Delta_{t,\lambda}(\theta) \tag{3}$$

For calculation of $D(\theta, n)$, total damage on the creature in a simulation, we add up all limb damages at the end with two considerations:

- As an incentive for bipedality, damages on the two most damaged ordinary limbs, i.e. limbs apart from head and torso, are ignored.
- Final head damage is multiplied by ten to be penalised more severely.

While we would be able to measure the internal damages (i.e. damages that spread from the joints) according to these formulae, we only focused on damage of colliding to the ground, as the perfect damage calculation tended to increase the simulation time considerably. Another possible improvement in realism would be to consider the surface area of the collision with the ground.

3.1.2. Locomotion and Perpendicular Area. If in a simulation the creature's head starts from the position (x_0, y_0, z_0) and after n time steps ends up at (x_n, y_n, z_n), and the movement of the head at the time step i in between is given by $(\Delta x_i, \Delta y_i, \Delta z_i)$, then the *perpendicular area* of the creature's movement is defined as:

$$A(n) = \begin{bmatrix} x_n - x_0 \\ y_n - y_0 \end{bmatrix} \cdot \sum_{i=0}^{n} z_i \cdot \begin{bmatrix} \Delta x_i \\ \Delta y_i \end{bmatrix} \quad (4)$$

which is guaranteed to be non-negative if for all i, $z_i \geq 0$. The movements which are parallel to the ground plane (xy plane) are measured and multiplied by the current height of head z_i; thus, the higher the head is, the better is the score, and here is the reward for *carrying* as well. It can be seen that $A(n)$ increases with successful locomotion in any direction, while vibration and back-and-forth motion do not receive a good score.

3.1.3. Basic definition of fitness. The final fitness will be proportionate directly to the perpendicular area of the movement, $A(n)$, and inversely to the total damage, $D(\theta, n)$:

$$f_{n,\theta} = \frac{A(n)}{1 + \mu D(\theta, n)} \quad (5)$$

in which the coefficient μ is to be fine-tuned by trial and error.

All creatures of a given generation are evaluated using the same fitness function f with the same n and θ; but in every new generation the fitness function will be slightly varied. We slightly increase n, number of time steps, to give a successful locomotor creature opportunity to continue moving and to gain more perpendicular area of movement. We also decrease the damage threshold θ according to a quadratic formula in every new generation to increase the sensitivity to damage.

3.1.4. Avoiding falling strategies. In the most of initial experiments we observed that the GA converged towards the creatures with strategies of controlled falling instead of successful locomotory skills. These creatures tended to get taller and taller by generation to compensate lack of their movement and gain a large perpendicular area only by falling. To deter this behaviour we tried a number of tricks and modifications.

The first tentative solution was increasing sensitivity to damage, which only deterred movement further but did not prevent falling strategies, only changed them to controlled falling strategies.

The final solution was using a threshold function on the height of head at each time step, whose threshold depends on the initial height of head, z_0. Formally, in (4) we substituted z_i by $T(z_i)$, defined as:

$$T(z_i) = \begin{cases} z_i & z_i \geq (1-\delta) z_0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

with $0 \leq \delta \ll 1$. We used $\delta = 0.2$.

3.2. Genetic Algorithm Process

In most of the experiments we used a population of size 1000 and aimed to run the GA for 100 generations. At generation zero, the population is generated randomly and then every creature is tested in a simulation to produce a fitness score. Then each creature's genome, generation and score are saved on the disk. This enables us to retrieve the GA if it halts at any point, and also makes it easier to use parallelism during processing.

We started with $n = 7500$ time steps for each simulation in the generation zero and increased it by 500 in every generation afterwards. As the fitness function was varied the convergence criterion could not be the convergence of the fitness score; therefore, we assumed that low variability in the morphology and behaviour of a generation is a sign of convergence.

As the simulations are the most time-consuming parts of the process, there are a few quality controls before the simulation test to avoid ill-formed configurations. For example, the creature should be checked for having just one head, as uniqueness of head is an essential assumption in the fitness function definition, and headlessness or polycephaly result in corrupted fitness scores.

After the population of each generation are fully produced, all the fitness scores are read from the disk and a selection method is used to choose two parents for each creature in the new generation. Then their genomes are read and a reproduction method, randomly chosen between a few alternatives like crossover, mutation, picking and grafting, is applied. Here we need the quality check again, as the offspring might be ill-formed despite its parents.

3.2.1. Selection method. Keeping adequate selection pressure can be vital for the success of any GA. According to Bäck and Hoffmeister [16] GA ideally should start with a low selection pressure to be *explorative* in the beginning and avoid entrapment in the local optima, then it should increase the selection pressure and become more *exploitative* to speed up the convergence. We used Boltzmann selection with *Standard Deviation Schedule* (SDS) which is an adaptive cooling schedule suggested by Mahnig and Mühlenbein [17]. The adequate initial temperature was found by trial and error.

As we encountered premature convergence, to reduce affects of allele loss we tried an elitist and preservative selection method in the GA, which keeps a few of the fittest creatures in the next generation

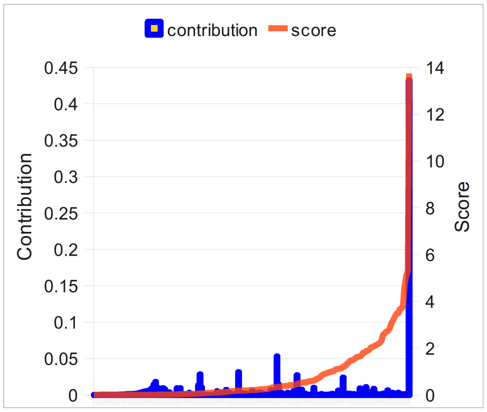

Figure 1. Genealogy in practice: relationship between scores of the generation-zero ancestors of F20 and their contribution

without reproduction. The size of this group is called the *elite size* which was often 25 in our experiments.

4. Results

4.1. Initial Experiments

We had a few unsuccessful attempts, where the GA reached a premature convergence to inelegant strategies. This could be seen as a tendency to generate homologous morphology and behaviour in all creatures of an early generation rather than convergence of their fitness score, as the GA was able to keep improving scores even within a limited and invariable gene pool.

As already mentioned in 3.1.4, we often observed falling strategies along with growing taller. In these morphologies, risk of head damages and dangers of falling behaviour were neutralised by growing an anchor limb that absorbed the impact.

4.1.1. Genealogy in practice. The main challenge in many GA applications is tackling premature convergence. In our case, with a varying fitness function and very high-dimensional search space, premature convergence is hard to trace, let alone to tackle. We observed the convergence in morphology and behaviour despite a steady growth in fitness scores. As we needed an objective measure for tracing convergence, we tried tracing the ancestors of creatures in generation zero to compare genome combination between two creatures of different generations, and if a creature of later generation has exactly similar ancestral combination as an earlier one, we consider it as a sign of convergence.

For example, in one experiment, in generation zero with population size 1000, the average fitness score was 0.3 with standard deviation 0.57, while the fittest of this generation, *F0*, scored 13.73 (the next best score was 5.27). If the distribution was normal this score would be extremely unlikely.

By applying genealogy to the same experiment we found out *F0* has 43% contribution in the fittest creature of 11^{th} generation *F11* and same contribution in the fittest creature of 20^{th} generation, *F20*. The next greatest contributor ancestor of generation zero, with 5% contribution in both *F11* and *F20*, had modest score of 0.33. Further analysis illustrates that the fitness of the contributors is independent of their share of contribution to the fittest creatures of the next generations, unless the fitness is distinctly higher than others (figure 1).

Although the average fitness score grows from 24.1 in generation 11 to 29.3 in generation 20, we could detect genetic convergence by using genealogy. In this condition we have an invariable gene pool in the later generations because of premature convergence, and mutations, not sexual reproductions any longer, are the main source of variability and improvement in fitness score.

4.1.2. Remedies for Premature Convergence. Premature convergence could be result of inadequate fitness function and/or out of tune GA parameters. Also, a very high dimensional search space could lead to the *curse of dimensionality*, where the random initialisation of points is insufficient to evenly cover the search space. As a result of the lack of variability, the GA is very unlikely to reach some particular points in few generations, so always ends up to the local optima.

Thus, we tried to reach a more even distribution of fitness scores in the generation zero to make the next generations more variable genetically. This is achieved by defining a minimum fitness score for a creature to be registered in the generation zero. We also severely increased the maximum allowed force that can be exerted to a joint by creatures, and obtained more versatile and active behaviours as result.

4.2. Biped creatures

We had two qualitatively successful experiments[1], which were both initialised with the creatures that passed the minimum score test described in 4.1.2. In experiment [a] the minimum score was 10 and in experiment [b] it was 0.5.

1. Videos of the evolved creatures can be viewed at http://tinyurl.com/bipedalism

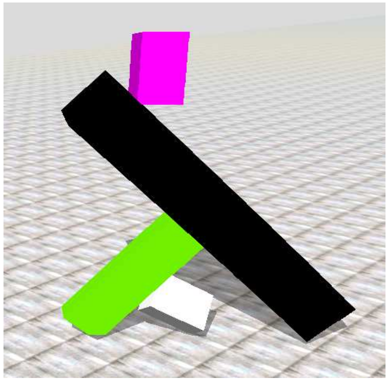

Figure 2. Biped morphology of fittest creature in generation 93 of [a]

4.2.1. The Galloper. The overwhelming emerged morphology and behaviour of [a] is a biped "galloping" creature whose fittest samples perform an excellent balance during hopping and travel in high speed, although the balance does not seem robust enough (figure 2). In one galloping cycle, the front foot is used for jumping and the back foot absorbs the falling impact. This causes damages which are illustrated by darker limb colours in the simulation, but this damage is ignored by the fitness function because of considerations mentioned in 3.1.1.

The upright morphology is apparent with the head (coloured magenta) on the top and the limbs (depicted in green or black) at the bottom. The torso (white block) seems to have become vestigial but it might help the balance. In the later generations, however, it is the main cause of damage penalties, since the feet are exempt from damage but the torso can get slight contacts with the ground, which are damaging in high sensitivities of later generations.

As getting score over 10 in generation zero by a completely random configuration has small probability (about 0.1%), to reduce the time required for experiment the population size was reduced from 1000 to 95, to fill which around $55.1k$ random creatures were generated, whose average fitness score was 0.4 and the standard deviation was 0.95, in which we see more versatility and slightly better average compared to the generation zero of the experiment described in 4.1.1 (average 0.3 and standard deviation 0.57), which is achieved by increasing maximum allowed force exerted to the joint.

Probably because of more uniform distribution of the scores, genealogy of the fittest creature in 30$^{\text{th}}$ generation (tracking back its ancestors to the generation zero, in figure 4) shows that there are 22 contributors, which is considerable compared to the population size 94, and the chance of contribution is distributed more evenly compared to the experiment described in 4.1.1. Figure 3 shows the evolving of the average and standard deviation over generations in [a]. The average starts at 14.7 with standard deviation 3.7, and gradually grows to 375.8 in the generation 91, then with the damage threshold falling very close to zero and increasing sensitivity the average drops rapidly to 1.66 in generation 100. Best average of [a] is 8.3 times greater than the best average of former unsuccessful attempts.

4.2.2. The Crawl-Stepper. In [b], the population size is 1000 with elite size of 25, and as the minimum score 0.5 was required, about $5.1k$ creatures were made to provide the required population in generation zero, coincidentally close to that needed in [a].

The experiment [b] resulted in one of the qualitatively best strategies of locomotion. Even though the morphology is not completely bipedal, the behaviour is very similar to walking, with an oddly enough stepping gait which takes advantage of only one limb. The creature shows dynamic balance, and while it is slower than the galloper, it is more robust. It stands on one horizontal limb (pelvis) and swings the upper body, which causes two sides of the pelvis periodically rise and fall in a stepping-like behaviour (figure 5).

The experiment went on for ten generations, where the convergence in behaviour and morphology was obvious. It had the fastest growth in the average scores amongst all the experiments with average of 286.2 in the generation 9, and the increase in the standard deviation is considerably less than the others (figure 6).

Despite the larger population, genealogy shows that the fittest creatures in the later generations all originated from a single ancestor in the generation zero, so the population size did not bring variability through the sexual reproduction and might only help by making the competition more severe and also increasing the number of mutations as the main drive of evolution. We also cannot claim that the population size compared to [a] increased the chance of better initialisation of genomes, as the number of tested creatures to provide generation zero are almost equal in both [a] and [b]. But the populations size definitely increased the number of simulation tests in each generation, so ten generations of [b] has almost as many tests as a hundred generations of [a] has.

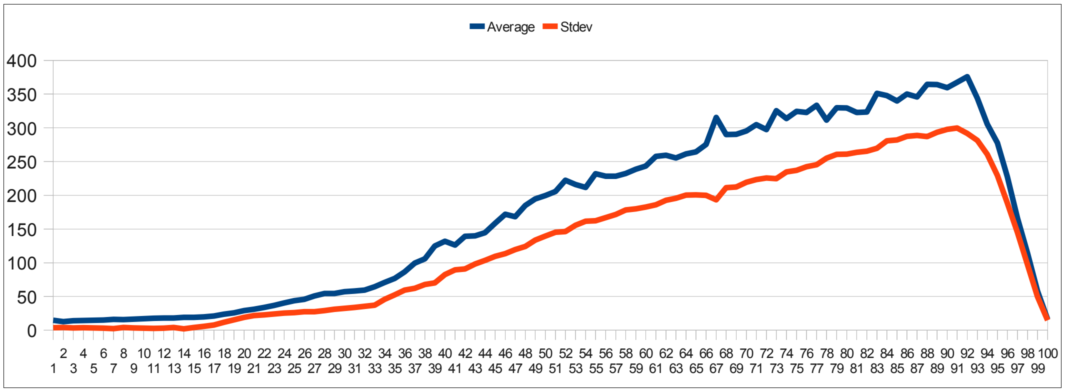

Figure 3. Average and standard deviation of scores in each generation in [a]

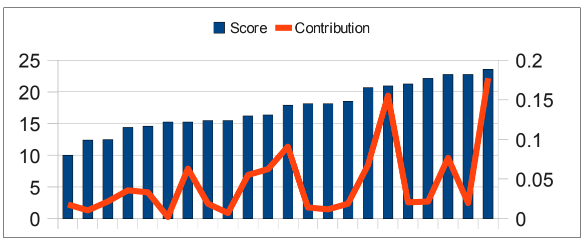

Figure 4. Scores of the creatures in generation zero which contribute to the genome of the fittest in the generation 30 in [a]

5. Future Work

The main focus of our work was on defining and tuning a fitness function. Elaborating evaluation of creatures still can be the main thrust of future work and lead to evolution of more realistic strategies.

A suggestion is measuring balance robustness of the creatures, which can involve simple alterations of the evolved neural network (e.g. by disconnecting some neurons or adding noise proportionate to the damage) or can involve external force, for example changing gravity, using uneven or shaky terrain, or exerting an external destabilising force, as done in [2].

Another is sexual and/or aesthetic selection, which may have played a complementary role in evolving hominids into upright locomotion as it is seen that the biped or quadruped creatures might be aesthetically and visually more pleasing to the human eyes. In simulation it can be done by human observers choosing the more visually attractive ones between the variants of the creatures to reproduce in the next generations (as Ray did in [18]).

The definition of damage still can be improved too: as Miconi points out [9, pages 43-44] default homogeneity of materials in the simulation worlds is a potentially limiting oversimplification. We can alter it, or even in a homogeneous world we can differentiate between the surface area of different types of collisions, that helps us to distinguish between a cutting edge and a patting touch. Also, internal damage due to force that spreads through joints can be considered.

To include all the needed characteristics a few consequent evolution stages can be planned, each with different fitness evaluations. For example, in addition to including penalties for damage from falling and collisions, and/or sexual or aesthetic selection, we could evolve merely for locomotion in the first generations and later introduce the reward for carrying abilities to better reflect evolutionary history of hominids.

6. Conclusion

We evolved bipedalism in the Sims's creatures by incorporating physical damage and incentives for upright locomotion. The reward for carrying is reflected in the components of the fitness function involving keeping the head up, limiting the number of limbs and making two limbs exempt from damage. The definition and tuning of the fitness function along with the fine tuning of the GA process were critical to the evolution of the two bipedal creature types obtained. The results support the plausibility of the argument that freeing the hands for carrying may have been a crucial factor in the evolution of bipedalism.

References

[1] K. Hirai, M. Hirose, Y. Hikawa, and T. Takenaka, "The development of Honda humanoid robot," in *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, vol. 2, May 1998, pp. 1321 –1326 vol.2.

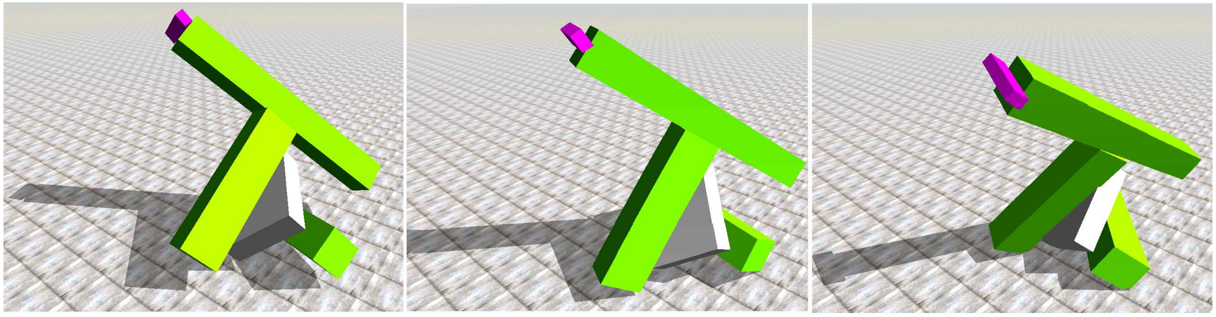

Figure 5. The fittest creature of [b] in simulation

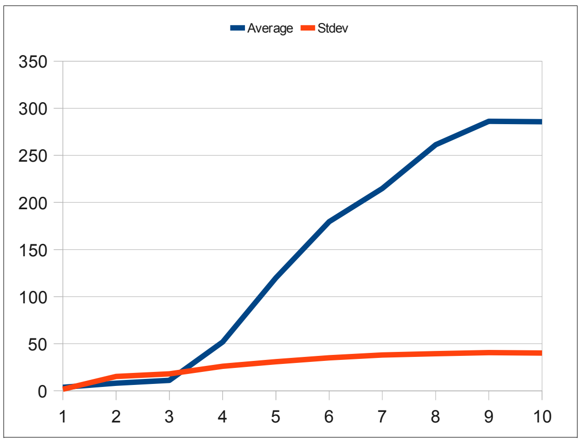

Figure 6. Average and standard deviation of scores of each generation in [b]

[2] E. D. Vaughan, E. D. Paolo, and I. R. Harvey, "The Evolution of Control and Adaptation in a 3D Powered Passive Dynamic Walker," in *In Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE9*. MIT Press, 2004, pp. 139–145.

[3] S. C. Morris, *Life's solution: inevitable humans in a lonely universe*. Cambridge: Cambridge University Press, 2003.

[4] C. O. Lovejoy, "Evolution of Human Walking," *Scientific American*, pp. 118–125, Nov. 1988.

[5] D. M. Bramble and D. E. Lieberman, "Endurance running and the evolution of *Homo*," *Nature*, pp. 345–352, Nov. 2004.

[6] K. Sims, "Evolving Virtual Creatures," in *Computer Graphics*, Jul. 1994, pp. 15–22.

[7] P. Prusinkiewicz and A. Lindenmayer, *The algorithmic beauty of plants*. New York, NY, USA: Springer-Verlag New York, Inc., 1996.

[8] K. Sims, "Evolving 3D Morphology and Behavior by Competition," in *Artificial Life IV*, 1994, pp. 28–39.

[9] T. Miconi, "The Road to Everywhere: Evolution, Complexity and Progress in Natural and Artificial Systems," Ph.D. dissertation, Univesity of Birmingham, Birmingham, UK, 2007.

[10] T. F. Novacheck, "The biomechanics of running," *Gait & Posture*, vol. 7, pp. 77–95, Jul. 1998.

[11] P. D. Moral and L. Micio, "Genealogies and Increasing Propagation of Chaos For Feynman-Kac and Genetic Models," *Annals of Applied Probability*, vol. 11, pp. 1166–1198, 2001.

[12] J. P. Rosca and D. H. Ballard, "Causality in genetic programming," in *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*. Morgan Kaufmann, 1995, pp. 256–263.

[13] J. R. Koza, "A hierarchical approach to learning the Boolean multiplexer function," in *Proceedings of Workshop on the Foundations of Genetic Algorithms and Classifier Systems*. Morgan Kaufmann, 1990.

[14] ——, "Architecture-Altering Operations for Evolving the Architecture of a Multi-Part Program in Genetic Programming," Stanford, CA, USA, Tech. Rep., 1994.

[15] T. Miconi, "Fitness transmission: a genealogic signature of adaptive evolution," in *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, S. Bullock, J. Noble, R. Watson, and M. A. Bedau, Eds. MIT Press, Cambridge, MA, 2008, pp. 404–411.

[16] T. Bäck and F. Hoffmeister, "Extended Selection Mechanisms in Genetic Algorithms," in *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann, 1991, pp. 92–99.

[17] T. Mahnig and H. Mühlenbein, "A new adaptive Boltzmann selection schedule SDS," in *Evolutionary Computation, 2001. Proceedings of the 2001 Congress on*, vol. 1, 2001, pp. 183 –190.

[18] T. S. Ray, "Aesthetically Evolved Virtual Pets," in *Artificial Life 7*, 2000, pp. 158–161.