

Complex Networks of Simple Neurons for Bipedal Locomotion

Brian Allen

Department of Computer Science
University of California, Los Angeles

vector@cs.ucla.edu

Petros Faloutsos

Department of Computer Science
University of California, Los Angeles

pfal@cs.ucla.edu

Abstract—Fluid bipedal locomotion remains a significant challenge for humanoid robotics. Recent bio-inspired approaches have made significant progress by using small numbers of tightly coupled neurons, called central pattern generators (CPGs). Our approach exchanges complexity of the neuron model for complexity of the network, gradually building a network of simple neurons capable of complex behaviors. We show this approach generates controllers *de novo* that are able to control 3D bipedal locomotion up to 10 meters. This result holds for robots with human-proportionate morphologies across 95% of normal human variation. The resulting networks are then examined to discover neural structures that arise unusually often, lending some insight into the workings of otherwise opaque controllers.

I. INTRODUCTION

Walking robots are long-standing goal for the field of robotics. Humanoid bipedal robots hold special sway. This is in-part due to the practical importance of their development. Anthropomorphic robots might be able to relieve humans from dangerous or difficult work while capitalizing on the wide variety of human-oriented services and systems available. In addition to the practical utility of humanoid morphologies, humans may also prefer interactions with machines that have familiar form and motion. With this perspective, we describe a methodology for the creation of humanoid robot controllers with the ultimate goal of reproducing the fluidity and function of human locomotion.

Our method is rooted in evolutionary robotics— a biologically inspired approach that loosely simulates the natural control systems, using artificial neural networks, and their design through evolutionary or learning processes [1]. Prior work on the biologically inspired control of legged robots has had two main focii. The first approach uses layered, feed-forward networks of simple neuron models [2], [3]. A second approach is the use of more complex and realistic individual neuron models in highly interconnected networks. The networks of the second approach resemble central pattern generators (CPGs). CPGs are neural structures observed in the spine of many vertebrates that are believed to be responsible for the generation of cyclical patterns, such as for locomotion control [4]. In part, the success of the CPG approach results from the effort applied to careful design and weight tuning [5]. Methodologies to automatically tune the interconnection weights of CPGs have had some success, both by evolutionary process [6] and by supervised learning [7].

In this contribution, we describe a method for the automatic synthesis of controllers for bipedal locomotion using

a third biologically inspired approach: gradually growing networks of simple neurons. Starting with a minimal network of simple neuron models, like the first approach described above, neural structures are gradually added through an evolutionary process. In forgoing the inherent cyclical dynamics of more complex neuron models, such as those typically used in CPG-based controllers, the approach relies instead on complex interconnections between the neurons. The choice to use a simple neuron model is deliberate and embodies a trade-off: the loss of inherently cyclical pattern generation in exchange for the ability to modify the behavior in small steps with topological changes. In comparison to previous work using CPGs for bipedal locomotion, our approach moves complexity out of the neuron model and into the network.


Our results show the *de novo* creation of neural networks capable of limited bipedal locomotion. Controllers are evolved using a range of human-scale morphologies. In addition, individual controllers can adapt, without additional training, to previously unseen morphologies. The networks are analyzed by searching for sub-graphs (motifs) that are found more often in successful walking controllers than would be expected by chance. This result provides some practical insight into the workings of the neural networks, slightly disturbing the veil that has historically separated successful “black-box” neurocontrollers from an understanding of their mechanism.

II. THE SIMULATED BIPEDAL ROBOT

The simulated robot is designed with a set of actuated degrees of freedom (DoF) found to be minimal while still allowing for an anthropomorphic walk. The masses and sizes of body links, as well as the range of motion of joints, are based on aggregate measurements of military personnel [8].

The angular position of each actuated joint is provided to the controller. Sensory data are scaled so that the full range of motion is reported as $[-1, 1]$. In addition, contact sensors on each foot indicate whether the foot is in current contact with the ground (-1) or not ($+1$), and sensors provide the distance of each foot from the current center of mass (CoM), as projected to the ground plane. The height of the waist segment and the linear velocity of the CoM projected to the ground plane are also provided as sensory input.

The controller specifies the target angle of each actuated joint. The corresponding actuator applies torque to drive the joint to the desired angle using PD control. The output signal of the controller ranges within $[0, 1]$, of which the central 0.8 range is linearly scaled to the full range of the joint,



(a) Anthropometric simulated biped robot. (b) Schematic diagram of robot.

Fig. 1. The robot has 15 degrees of freedom (DoF) in total. Seven of those are actively controlled by the neurocontroller. The remaining DoFs act under PD control and maintain a constant target position defined by the standing-upright pose.

leaving 0.1 units at each extreme clamped to 0.0 and 1.0, respectively.

Joint, axis	Range		Controlled
	High	Low	
Spine, transverse	$-\pi/12$	$\pi/12$	Yes
Spine, coronal	$-\pi/16$	$\pi/16$	No
Spine, sagittal	$-\pi/16$	$\pi/16$	No
Hip, transverse	0	0	No
Hip, coronal	$-\pi/48$	$\pi/24$	Yes
Hip, sagittal	$-\pi/16$	$\pi/16$	Yes
Knee	0	$3\pi/4$	Yes
Ankle, transverse	$-\pi/12$	$\pi/12$	No
Ankle, coronal	$-\pi/12$	$\pi/12$	No
Ankle, sagittal	$-\pi/4$	$\pi/4$	No

TABLE I

THE ALLOWED RANGES FOR EACH JOINT OF THE AGENT. THE CONTROL COLUMN INDICATES IF THE JOINT'S TARGET ANGLE IS SET BY THE CONTROLLER. IF NOT, THE JOINT APPLIES TORQUE TO CENTER ITSELF IN ITS RANGE.

III. NEURAL CONTROLLER

A. Neural Network

The neural networks used as controllers are composed of a set of neurons, each with a single scalar activation level y , and a set of weighted, directed connections between the neurons.

Each neuron's activation level is updated by linearly combining its incoming connections' activations and applying the activation function, $\sigma(x)$, to that sum to determine its activation level for the next cycle:

$$y_i = \sigma\left(\sum_{j=1}^N w_{ji}y_j\right), \quad \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (1)$$

where σ is the sigmoid function and w_{ji} the weight of the connection from the j th neuron to the i th. Neural updates occur every 0.07 seconds of simulated time. This delay was chosen to be within the observed range of spinal reflex response times in humans [9].

To achieve state and time dependent behavior, the neural network must rely on cycles in its topology, in a manner analogous to cross-linked XOR gates forming an electronic flip-flop. Note that in the sigmoidal neurons used here, such cycles are the only way a network can store state internally. To manage complex timing-based behavior, as is needed for a motion controller, a network of simple neurons requires a complex topology of inter-neuron connections.

B. Bilateral Symmetry

Bilateral symmetry of control was enforced by using two identical and independently evaluated neural networks, each controlling one side of the robot. The outputs of each network drive the actuated joints on that side. Since the central waist joint is not associated with either side, it is controlled by the average of the activations of the waist-output neuron on each network. The inputs of each network are likewise set from per-side information. Although the topologies of the two networks are identical, the internal activation states are independent. In fact, since asymmetric motion is generally required to initiate walking from a standing pose, the internal states of the two control networks tend to diverge rapidly during locomotion.

Each neural network has two bias nodes. The first bias node has a constant activation of 1.0 and is the same for both the right and left networks. The second bias node has an activation of 1.0 for the right-side network and -1.0 for the left. This per-side bias allows for asymmetrical behavior even when sensory inputs are the same from both side. Two tactile foot-contact sensor inputs are connected as same-side and opposite-side.

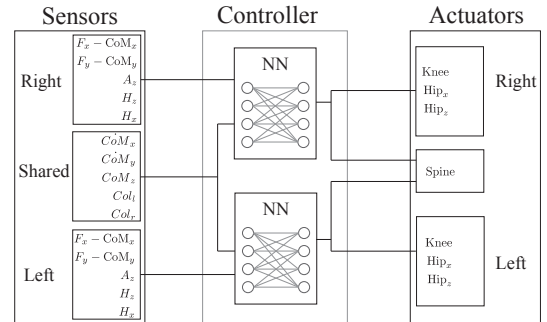


Fig. 2. The controller is composed of two structurally identical neural networks (NN), each accepts input signals from the shared and same-side sensors. Likewise, each neural network drives the actuators of one side and contributes to the control of the shared actuator.

IV. GENETIC ALGORITHM

The network topology and connection weights of the neural controllers are determined by a genetic algorithm. Our approach is heavily based on neuroevolution by augmenting topology (NEAT) [10]. As in NEAT, the evolved genome is represented as a set of connections between neurons, with real-valued weights that correspond directly to the neural connection weights w_{ij} in equation 1. The population size was 512 individuals. In each generation, every genome was

used to create two identical neural networks (right and left), and the resulting controller was evaluated according to the objective function (see section IV-D). Individual genomes were divided into species based on network similarity (see section IV-C), and the best-performing 25% of each population were preserved, with the remaining population filled by mating and mutating the best-performers to form the next generation.

A. Scalar and Topological Mutations

Mutations to the directed-graph genotype are of two kinds: structural and non-structural. The latter alter the weight of a connection or a parameter of a neuron. Small alterations ($\mu = 0$, $\sigma^2 = 0.1$) are most likely, though occasionally entirely new values are chosen (at random uniformly in the range $(-0.1, 0.1)$) as replacements. In an effort to bias mutational changes toward newer structures, parameters that are relatively new features in the genome are more likely to be replaced than older ones.

Three types of mutation affect the neural topology. The first splits an existing edge into two edges and a node. The second introduces a directed edge between two previously unconnected nodes, or between a node and itself. The third removes an edge, and if the removal results in interior nodes with no remaining connections, the nodes are removed as well.

B. Cross-Over

The key innovation of NEAT is the introduction of “historical markers” to label each neuron and connection. During asexual reproduction, these markers are preserved and passed to the offspring. During sexual reproduction, these historical markers are used to determine genetic homology. The guiding assumption is that genes with the same historical origin (and therefore the same historical markers) will perform the same function in the phenotype. Although this is, in essence, an ad hoc approach with a key assumption that can, and may often, fail, it has been shown to be among the most successful machine learning approaches for simple benchmark control problems such as double-pole balancing and predator-prey simulations [11].

C. Speciation

In addition to enabling a more productive cross-over operation, NEAT historical markers can also be used to estimate the chance of mating success between individuals. A distance metric groups the individuals of a population into species, allowing sexual reproduction within a species containing similar individuals. Such speciation greatly improves the likelihood that the next generation will be viable. The distance metric is defined as

$$\delta = \frac{G}{N^k} + cW, \quad (2)$$

where G is the number of genes without a corresponding historical marker in the genes of the other parent. The constant c is a normalization factor based on the magnitude

of connection weights. N is the number of genes in the larger genome, and k is a term allowing a scaling of the effect of normalization based on the number of genes N . Note that previous implementations of the NEAT algorithm have used $k = 0$, while our implementation uses $k = 1$. The effect of $k = 1$ is to measure genomic distance by the ratio of differing genes to the total number of genes, in contrast with $k = 0$, which considers the absolute number. Comparing the ratio, rather than absolute number better supports genomic distance calculations for the large networks needed for bipedal locomotion. This modification was found useful to ensure proper speciation for larger neural networks.

New *structural* innovations will likely differ significantly (as measured by δ) from existing genomes, meriting classification in a new species. Individuals only compete directly with other members of their own species, providing protection to new innovations until they have time to optimize their structure. This process is further assisted by giving new species an artificial fitness bonus of 20% for their first few generations.

D. Objective Function

Designing an objective function measure that accurately quantifies the quality of an arbitrary attempt at bipedal walking is a challenging notion. One might reasonably include a host of previously suggested factors, such as minimal energy use [12], or similarity to recorded human motions. Our approach is to instead use a very simple measure of the fitness of a walk. In doing so, we rely on the anthropometric body and torque limits to bias the search toward a human-like walk.

Our objective function is

$$f_{walk} = k_d \max(\|\text{proj}_{\vec{j}} \vec{d}\|, \epsilon), \quad (3)$$

where \vec{d} is the vector from the starting position to the hindmost foot, and \vec{j} is the unit vector in the direction the character is to walk. k_d is a constant scaling factor, and ϵ is a small positive value. We also experimented with more complex fitness functions for locomotion, however 3 was equally effective and used for the results presented here.

1) *Early Termination*: Identifying criteria for early termination of a trial is a useful way to improve the overall speed of the evolutionary process, since the evaluation of unpromising individuals is stopped as soon as one of the conditions is met.

Waist height: If the z -coordinate of the waist segment’s center of mass falls below a minimum height (50% of the waist segment’s starting height), the simulation is terminated and the genome’s final fitness is that computed in the previous time-step.

Instability: If, during the course of the physical simulation, significant numerical instability or joint divergence is detected, the simulation is terminated and the genome’s final fitness is set to the minimal allowed value, ϵ .

2) *Support Harness*: A simulated harness provides lateral and vertical and torsional support to the character. The strength of the harness is gradually diminishes from full-strength to exerting no force over the first 150 generations according to $h_s = 1 - \sin(\frac{\pi}{2} \frac{g}{150})$ where g is the current generation. In addition to the stabilizing support, the harness exerts a linear force in the direction of desired motion of $h_s m(v_d - C\dot{O}M_y)$ where m is the total scalar mass of the robot. The desired CoM velocity v_d is intended as a natural walking speed for the robot. Since we examine a variety of robots with different dimensions and masses, the natural walking speed must be generalized to physical properties. The Froude ratio Fr provides a simple method to estimate locomotion speeds based on the inverted pendulum model [13]. Using this method, the desired CoM velocity is $v_d = \sqrt{glFr}$, where l is the leg length of the robot and g the acceleration due to gravity of $9.81m/s^2$. $Fr = \frac{1}{4}$ is used to match a typical human-style walking gait [14].

V. RESULTS

The approach presented generates neural networks capable of controlling the biped walking distances of up to 10 meters. Fully stable controllers, able to walk arbitrarily long distances, were not produced. The authors suggest that the subjective appearance of the resulting controllers is interestingly organic and, considering that no motion trajectories were provided to the system, surprisingly human-like. Example joint angles for a typical successful walk controller are shown in figure 3. This particular individual is of medium female height, weight and hip-width.

Although evolutionary processes are stochastic by nature, our approach is reliably able to find controllers capable of walking short distances. Over the entire range of human body shapes evaluated at walking, 93% of the runs found controllers capable of walking upright for at least two meters (28 out of 30 runs). The failures were for the heaviest (95th percentile weight) and shortest (5th percentile height).




Fig. 3. Joint angles observed for a robot with media female morphology during the initiation phase and through several steps of a walk cycle.

A. Anthropometric Variation

Since we are partly concerned with controllers that provide human-like motion, we use robot morphologies whose

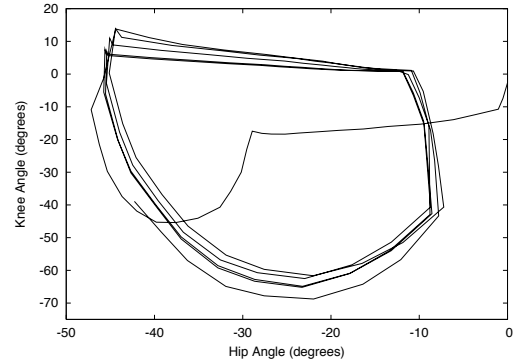


Fig. 4. Hip/knee joint cycle in walking, including the gait-initiation phase.

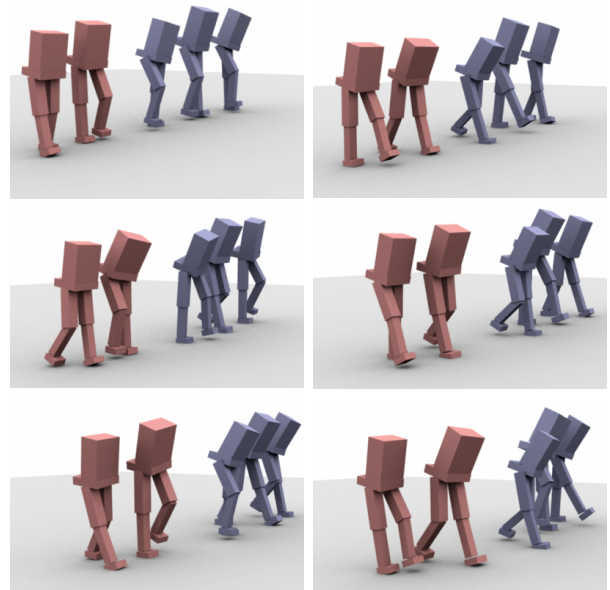


Fig. 5. Five characters starting to walk from a line-up. The two left-most (red) characters are 5th and 50th percentile height, weight and BIB females. The three remaining characters (blue) are males, with the nearest character 5th percentile height and 95th percentile weight and BIB, the middle blue character 95th height but 5th percentile weight and BIB, and finally the farthest character 95th percentile in each height, weight and BIB. Each of these characters evolved the behavior shown in a single evolutionary trial.

physical parameters mimic human dimensions. In addition, the generality of the approach is an important measure of success. To test the generality of our method in finding human-style gaits, we use a sample of humanoid character models from the normal human range of height, weight and bi-iliac (hip) breadth (BIB) covering the 5th to 95th percentile of each men and women.

Body segment sizes and weights are scaled linearly according to the three measures (height, weight and BIB) obtained from anthropometric data aggregated from over 30,000 individuals [8]. The relative proportions of segment sizes and weights are fixed and listed in Table II.

Neural networks evolved controlling a single robot morphology (height, weight and BIB). We then applied that network to the control of different robot morphologies to

Body Segment	Height Proportion	Mass Proportion
Trunk	31.2%	51.6%
Waist	6.2%	10.3%
Thigh (2)	31.0%	19.3%
Shank (2)	25.3%	9.3%
Foot (2)	6.2%	9.3%

TABLE II

THE ROBOT'S MORPHOLOGICAL PROPORTIONS REMAINED CONSTANT AND ANTHROPOMETRIC. THE MASS PROPORTION OF BILATERALLY REPEATED SEGMENTS SHOWS THE SUMMED MASS PROPORTION OF BOTH SIDES.





Fig. 6. This graph shows the effect of varying the physical model on the fitness of a controller. The network evolved using a single robot morphology with median female proportions. The dashed curve shows the distribution of the population of human female heights, approximated as a Gaussian distribution [8]. The dashed horizontal line is the fitness of the neural network when controlling the robot used for the network's evolution.


test the degree of generalization of the controller. With robot morphologies within around one standard deviation of normal human variation, the network is able to take a few steps, traveling about half the distance it reaches when applied to the morphology it evolved with. The results of variation in height across a human-like range are shown in Figure 6. That the network is able to adapt to robots with large changes in simulated morphology hints that our method might be robust to differences between simulation and a physical robot, potentially reducing the problem of the “reality gap” [15].

B. Patterns in the Network

The controller networks resulting from the evolutionary search are generated through a random process tempered by selection. As described, the structural mutations are fully random— any potential edge is as likely as any other to be added, and any existing edge is as likely as any other to be split with a new neuron. Given that all evolutionary changes to the neural network's structure are random, one might expect for the resulting networks to resemble random networks. However, the networks evolved for locomotion contain recurring patterns of interconnections that are significantly non-random. Recent work in the biological sciences has proposed methods for finding such patterns, or *motifs*, and determining the significance of their presence in a collection of networks.



(a) $p = 0.001$ (b) $p = 0.004$ (c) $p = 0.006$ (d) $p = 0.019$



(e) $p = 0.028$ (f) $p = 0.032$ (g) $p = 0.04$ (h) $p = 0.044$

Fig. 7. A selection of some of the motifs common to successful controller networks. Each subgraph is labelled with its p -value, that is, the probability that the frequency of that motif in a randomized network is at least the frequency of the motif in the controller network.

To identify significant motifs, we use the MAVisto [16] software to compare the instances of each 3-node and 4-node motifs in the best-performing network of the 500th generation for four different locomotion searches. Each search uses a robot with median height, weight and BIB for a human male. All four controllers are capable of walking at least five meters. One of the four networks is reproduced in Figure 8(e). Statistical significance of motifs is determined for each network individually by comparing the motif frequencies found in the controller network to frequencies found in one thousand randomly generated networks with the same distribution of vertex degrees. Self-loops were not included. Figure 7 shows a selection of motifs that are significantly ($p < 0.05$) more frequent in the evolved networks than in random networks. Subfigures 7(g) and 7(h) show a span of three neurons in mutual feedback, contributing a single output. Variations of cycles dominate. This is notably in contrast to similar analysis conducted with general biological neural networks [17], but in agreement with the general structure of CPGs.

Although not explored here, such motifs may have potential to improve the evolutionary search process. Ideally, the random changes that occur as part of the mutation process would not be truly random, but instead would be biased toward network topologies likely to exhibit the desired behavior. Identifying particularly common or rare motifs represents a first step in investigating this hypothesis.

C. Alternative Strategies

An evolutionary approach is rooted in specifying only a generalized description of the desired behavior, and the results can occasionally be unexpected or use unusual strategies. Some evolutionary runs produced successful motion strategies that differed from the desired human-like locomotion. Figure 9 shows a time-series sequence of one such alternative strategy, a one-legged hopping motion. Note that ankle target angles were not controlled by the neural network and this motion results from coordinated use of the knee and hip joints. Two additional strategies are shown in the accompanying video.

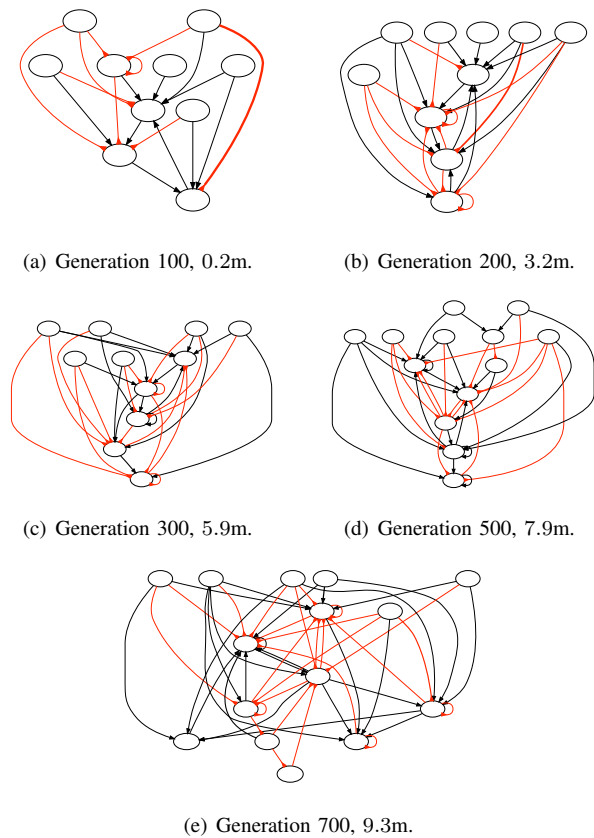


Fig. 8. Successive networks in a single evolutionary run show a gradually increasing topology. The thickness of the connecting arrow corresponds to the absolute value of the connection weight, with excitatory (positive-valued) connections in black and inhibitory (negative-valued) connections in red with inverted arrow-heads. Each network is labeled with the generation of the network and the distance travelled.

VI. CONCLUSION

This work presents evolved neural network controllers that show smooth motion, without the stiffness and phase artifacts generally associated with methods based on a fixed set of phases or states. The controllers are closed-loop, using proprioceptive and tactile sensors to maintain balance for limited periods.

The contributions of this work are (1) the use of NEAT for the neuroevolution of biped controllers, (2) demonstrating the generality of the method to a range of human-like morphologies, (3) demonstrating the adaptability of individual controllers to different morphologies, and (4) illuminating the common structure of the evolved networks using motif analysis.

REFERENCES

- [1] S. Nolfi and D. Floreano, "Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines," *MIT Press*, 2000.
- [2] J. C. Bongard and C. Paul, "Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution," in *From Animals to Animats: The Sixth International Conference on the Simulation of Adaptive Behaviour*. MIT Press, 2000, pp. 420–429.




Fig. 9. Some evolved controllers use alternative strategies to maximize the objective function, such as this one-legged hopper.

- [3] C. Paul and J. C. Bongard, "The road less travelled: Morphology in the optimization of biped robot locomotion," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 226–232, 2001.
- [4] P. M. Dimitrijevic MR, Gerasimenko Y, "Evidence for a spinal central pattern generator in humans," *Ann N Y Acad Sci*, vol. 860, pp. 360–76, Nov 1998.
- [5] G. Taga, Y. Miyake, Y. Yamaguchi, and H. Shimizu, "Generation and coordination of bipedal locomotion through global entrainment," *Autonomous Decentralized Systems, 1993. Proceedings. ISADS 93., International Symposium on*, pp. 199 – 205, Jan 1993.
- [6] T. Reil and P. Husbands, "Evolution of central pattern generators for bipedal walking in a real-time physics environment," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 159–168, 2002.
- [7] L. Righetti and A. Ijspeert, "Programmable central pattern generators: an application to biped locomotion control," *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pp. 1585 – 1590, Apr 2006.
- [8] Department of Defense, United States of America, *Anthropometry of U.S. Military Personnel (DOD-HDBK-743A)*. Department of Defense, United States of America, 1991.
- [9] D. Dewhurst, "Neuromuscular control system," *IEEE Transactions on Biomedical Engineering*, vol. 3, pp. 167–171, 1967.
- [10] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [11] —, "Competitive coevolution through evolutionary complexification," *Journal of Artificial Intelligence Research*, no. 21, pp. 63–100, 2004.
- [12] C. Chevallereau, A. M. Formal'sky, and B. Perrin, "Low energy cost reference trajectories for a biped robot," in *Proceedings of the IEEE International Conference on Robotics and Automation*. Piscataway: IEEE Computer Society, 1998, pp. 1398–1404.
- [13] C. Vaughan and M. Oamalley, "Froude and the contribution of naval architecture to our understanding of bipedal locomotion," *Gait & Posture*, vol. 21, no. 3, pp. 350–362, Apr 2005.
- [14] A. Biknevicius and S. Reilly, "Correlation of symmetrical gaits and whole body mechanics: debunking myths in locomotor biodynamics," *Journal of Experimental Zoology Part A: Comparative Experimental Biology*, vol. 305A, no. 11, pp. 923–934, 2006.
- [15] N. Jakobi, P. Husb, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life*. Springer-Verlag, 1995, pp. 704–720.
- [16] F. Schreiber and H. Schwobbermeyer, "Mavisto: a tool for the exploration of network motifs," *Bioinformatics*, Jan 2005.
- [17] R. Milo, S. Shen-Orr, S. Itzkovitz, and N. Kashtan, "Network motifs: simple building blocks of complex networks," *Science*, Jan 2002.