A STEP TOWARD EVOLVING BIPED WALKING BEHAVIOR THROUGH INDIRECT ENCODING

by

RANDAL S. OLSON

A thesis submitted in partial fulfillment of the requirements
for the Honors in the Major Program in Computer Science
in the College of Electrical Engineering and Computer Science
and in The Burnett Honors College
at the University of Central Florida
Orlando, Florida

Spring Term 2010

Thesis Chair: Dr. Kenneth O. Stanley

Abstract

Teaching simulated biped robots to walk is a popular problem in machine learning. However, until this thesis, evolving a biped controller has not been attempted through an *indirect encoding*, i.e. a compressed representation of the solution, despite the fact that natural bipeds such as humans evolved through such an indirect encoding (i.e. DNA). Thus the promise for indirect encoding is to evolve gaits that rival those seen in nature. In this thesis, an indirect encoding called *HyperNEAT* evolves a controller for a biped robot in a computer simulation. To most effectively explore the deceptive behavior space of biped walkers, *novelty search* is applied as a fitness metric. The result is that although the indirect encoding can evolve a stable bipedal gait, the overall neural architecture is brittle to small mutations. This result suggests that some capabilities might be necessary to include beyond indirect encoding, such as lifetime adaptation. Thus this thesis provides fresh insight into the requisite ingredients for the eventual achievement of fluid bipedal walking through artificial evolution.

Acknowledgments

First and foremost, I would like to thank my mother, Linda Olson-Avery, for her loving care and endless support through college. Her words of encouragement are what inspired me to always keep my chin up through the hard times and give the best that I have to offer.

Special thanks to my Thesis Chair, Dr. Kenneth Stanley, who took a scruffy Computer Science student with an interest in AI under his wing and spent the countless hours teaching, and sometimes re-teaching, me everything I needed to know to complete this thesis. With his guidance, support and encouragement, I gained the confidence to complete this tremendous undertaking and learned more than I ever could have imagined.

Joel Lehman in the University of Central Florida EPlex Lab, whose previous research with NEAT in the biped walker domain played a tremendous role in my success with this thesis. Joel's assistance and advice as I was initially getting my experiment off the ground was invaluable.

To everyone else who helped me on this thesis along the way, I give my sincere gratitude.

Contents

List of Figures

List of Tables

1 Introduction

For the past two decades, digital artists and engineers have sought to understand and reproduce human bipedal locomotion. Digital artists have elegantly reproduced human bipedal locomotion in films and video games while engineers have attempted to delicately engineer bipedal robots to walk like humans. While both enterprises have met with varying levels of success, their methods suffer from the need for fine-tuning the model to perform specific tasks. If a new task is required of the bipedal model, it must be modified and again fine-tuned by experts to perform the new task. In contrast, humans are capable of learning and adapting to new bipedal locomotion tasks as they are encountered with relative ease (e.g. walking up and down a flight of stairs for the first time). In principle, because dynamic bipedal control evolved in humans, computers may be capable of such a task as well. Thus the goal of this thesis is to take a step towards evolving a robust bipedal control system that can learn to walk a bipedal model as far as possible, without providing any prior knowledge on how to walk or even what walking is.

Solving this problem is important because it would open up many possible applications. In the field of robotics, evolved bipedal control systems could replace painstaking fine-tuning by robotics experts and provide a more dynamic bipedal control system for a bipedal robot. Within video games, evolved bipedal models could provide a more dynamic and realistic experience for viewers because the model is capable of reacting dynamically to new situations, instead of reacting in a pre-programmed fashion that looks stereotypical. Furthermore, in animated film, evolved bipedal models could save digital artists significant time by providing a basic walking model that can be fine-tuned with little effort to walk with a desired gait. Because of its many possible applications, the biped walker problem is a popular area of research in computer science and, more specifically, machine learning.

1.1 The Biped Walker Problem

Ultimately, the goal of the biped walker problem is to create a controller for a biped model that is capable of walking much like a human biped. That is, the problem seeks to produce a biped walker controller which not only walks with a stable, oscillatory gait, but is also capable of adapting to different terrains and biped models with relative ease. That way, if the biped walker controller encounters a new terrain (e.g. such as encountering stairs after walking on a flat surface) or a new model (e.g. the biped model's foot is damaged), the controller is robust enough to adapt to these changes and continue walking with a stable, oscillatory gait. Thus an effective controller for a biped walker is not only capable of balancing and oscillation, but also of adaptation.

1.2 Contributions

The set of experiments in this thesis explore for the first time the potential for *indirect encoding*, which means a compressed representation of the solution controller, to evolve effective biped walking policies. Experimental results confirm that indirect encoding in fact can evolve oscillatory gaits and the ability to balance for several meters; however, the bipeds still suggest room for improvement, implying that a further ingredient may ultimately need to be added to such systems to produce genuinely fluid motion and seamless balance. Thus the major contributions of this thesis are to lay the groundwork for further work in indirectly evolving bipedal walking, and to suggest the need for an added capacity for refinement, which may be made possible by evolving *adaptive* neural networks in the future.

1.3 Outline

The thesis commences with a brief review of NEAT, HyperNEAT and CTRNNs. Introduced in Section 3 are the Substrates, Functional Modularity equations and the Novelty Search fitness metric as they are implemented in the Biped Walker experiment. The Biped Walker experiment is described in Section 4. In Section 5, the results of the Biped Walker experiment are provided and analyzed. Discussed in Section 6 are the insights gained from this thesis and the proposed course of research stemming from these insights.

2 Background

This thesis builds on a representation called *compositional pattern producing networks* (CPPNs), which have been found to generate patterns with geometric regularities (Stanley 2007). The further idea that CPPNs can encode the connectivity of neural networks is the basis of the HyperNEAT algorithm (Stanley et al. 2009), which evolves the bipeds in this thesis. By utilizing CPPNs to express the connectivity of a neural network, HyperNEAT is capable of expressing neural networks with complex connectivity patterns that are sensitive to the geometry of the problem (Clune et al. 2009b; Gauci and Stanley 2008, 2007; Stanley et al. 2009). The interesting potential of CPPNs is that the inherent geometry in the biped walker problem may be possible to exploit. HyperNEAT has also been shown to perform well on similar problems with high regularity, such as quadruped walking (Clune et al. 2008). Both NEAT and HyperNEAT are reviewed in this section.

2.1 Biped Walking in Neuroevolution

Evolutionary algorithms and neuroevolution techniques have been applied to the biped walker problem domain with some success in the past (Allen and Faloutsos 2009; Hein et al. 2008; Reil and Husbands 2002; Van de Panne and Lamouret 1995). However, these attempts utilized more simple encodings that did not aim to learn regularities. Instead, researchers traditionally build in some knowledge of the domain from the start, such as enforced symmetry or oscillation. These constraints make evolution easier, but they also limit the creativity of its results. As of this writing, HyperNEAT has not yet been successfully applied in the biped walker problem domain. Furthermore, the hope is that HyperNEAT's ability to learn from geometry will help it evolve biped walking even without a priori constraints provided by the experimenter.

Finally, *continuous time recurrent neural networks* (CTRNNs) have proven useful in the evolution of oscillatory behaviors that are common in bipedal controllers (McHale and Husbands 2004), so HyperNEAT will be augmented in this thesis to evolve this type of network. Furthermore, novelty search (Lehman and Stanley 2008), which has been shown to avoid deception in some problems, will provide an alternative to the fitness function because the biped walking problem is notoriously deceptive (Allen and Faloutsos 2009).

NEAT and HyperNEAT are reviewed next.

2.2 NEAT

HyperNEAT is an extension of the Neuroevolution of Augmenting Topologies (NEAT) algorithm for evolving ANNs, which has performed well in a number of control and decision-making problems (Stanley and Miikkulainen 2002; Stanley et al. 2005; Lehman and Stanley 2008). NEAT begins with a population of simple ANNs and complexifies them over a series of evolutionary generations by adding and removing nodes and connections via mutation. By evolving networks in this fashion, NEAT is capable of generating increasingly complex ANNs until an acceptable level of complexity is discovered for the problem domain. Therefore, NEAT generally evolves the simplest ANNs necessary to solve the problem. In addition, high-level functionality discovered early in the evolutionary process can be maintained and elaborated upon in later generations. Stanley and Miikkulainen (2002) provides a more detailed description of the NEAT algorithm and how it operates.

2.3 HyperNEAT

The primary appeal of HyperNEAT for this experiment is its use of an *indirect* encoding (Bentley and Kumar 1999; Bongard 2002; Hornby and Pollack 2002) of ANNs, which is different from many neuroevolution techniques, including NEAT. Instead of *directly* encoding the ANN, in which each part of the solution's representation corresponds to a single component in the final ANN (Hornby and Pollack 2001; Stanley and Miikkulainen 2003; Reisinger et al. 2005), HyperNEAT concisely describes the connectivity of the final ANN with a simple description; the description of the ANN is often significantly smaller than the final ANN described (Stanley et al. 2009; Gauci and Stanley 2010). By describing the final ANN in such a fashion, high-level functionality, once discovered, can easily be re-used in the final ANN without having to re-discover the functionality in another area (Stanley 2007). In the biped walker domain, this capability can potentially help: the repetitive, oscillatory behavior of walking can be discovered for one leg and mirrored onto the other leg of the walker by simply copying the description over for the other leg. HyperNEAT, which is reviewed in this section, has succeeded in numerous difficult domains in which discovering regularities in connectivity patterns is known to be important (Gauci and Stanley 2008; Clune et al. 2009b,a; Stanley et al. 2009). A more detailed description of HyperNEAT is provided in Stanley et al. (2009) and Gauci and Stanley (2010).

In the HyperNEAT algorithm, NEAT still evolves a simple ANN; however, this ANN is no longer the final ANN for the solution, but instead the indirect encoding of the final ANN. This special ANN that encodes a second ANN is called a *compositional pattern producing network* (CPPN; Stanley 2007). CPPNs encode compositions of functions, where each function (which is the activation function of each node) is loosely related to a useful geometric regularity. Of special interest in the biped walking domain, CPPNs are capable of using a Gaussian function to produce symmetry in the final ANN. The appeal of encoding the final ANN by this method is that patterns of connections can be represented as a network of simple functions. In short, NEAT still evolves an ANN, but this ANN (called a CPPN) *encodes* the final network instead of being the final network itself.

The *substrate* is a set of nodes with a pre-defined geometry whose connectivity is decided by the CPPN (Stanley et al. 2009). By providing a fixed-topology ANN as the substrate, the experimenter can provide knowledge about the geometry of the problem to HyperNEAT. HyperNEAT is then able to exploit the geometry of the problem, which was previously not possible with directly-encoded ANNs. A detailed description of the substrates used in the experiments in this thesis is provided in the next section.

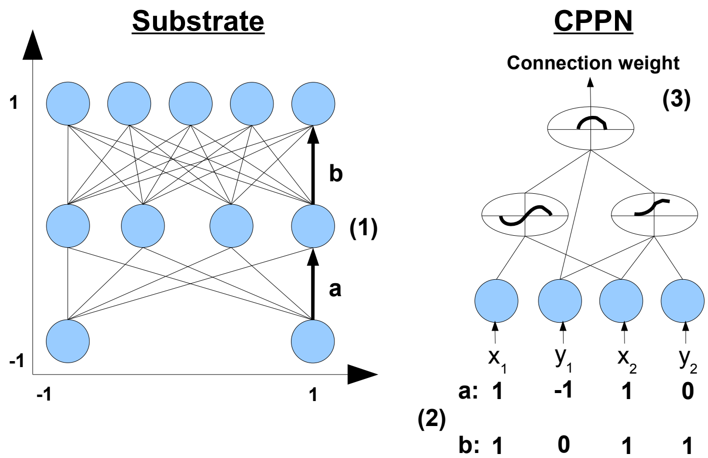

Figure 1: **CPPN-based Encoding of a Substrate.** The process by which HyperNEAT generates a substrate is demonstrated. The CPPN (3) generates the connection weights on the substrate (1).

The substrate connectivity is generated by first querying every possible connection between nodes in the substrate (Figure 1, Step 1): the coordinates of the two connecting nodes are provided as inputs to the CPPN (Figure 1, Step 2). Once activated, the CPPN outputs a weight, which is assigned as the weight of the connection between the two nodes (Figure 1, Step 3). If the generated weight is too low, the connection is not expressed. For example, to generate the connection weight between the nodes at coordinates (x_1, y_1) and (x_2, y_2), the coordinates x_1, y_1, x_2, y_2 are given as input to the CPPN. The CPPN is then activated and the resulting output is assigned as the connection weight between the two nodes, if the generated weight exceeds the minimum connection threshold. By performing this procedure for every possible connection within the substrate, the CPPN generates a connectivity pattern over the substrate. Figure 1 visually demonstrates this process.

2.4 Continuous Time Recurrent Neural Networks

In this thesis HyperNEAT is modified for the first time to encode *continuous time recurrent neural networks* (CTRNNs). CTRNNs have often been applied in the biped walker domain in the past because they maintain stability even with large, sudden changes in their inputs (McHale and Husbands 2004). This stability is accomplished by gradually integrating the output of every node over a given amount of time throughout the simulation, called the *time constant* for that node. The significant change in HyperNEAT to encode CTRNNs is that CPPNs also generate the time constants and biases for every node in the substrate. The CPPN generates the time constant and bias for a node similarly to how it generates the connection weight between two nodes, except instead of supplying x_2, y_2 as the input for the second node, the values 0.0, 0.0 are supplied instead. That way, CPPN can return node-centric as opposed to connection-centric values.

3 Methodology

This section describes the methods, techniques and algorithms that are introduced in the biped walker experiment.

3.1 Biped Substrates

An important factor in creating experiments with HyperNEAT is the design of the substrate. Two substrates were designed for use with the biped walker experiment: a two-dimensional substrate and a three-dimensional substrate. Both substrates were designed with the biped walker model geometry in mind to best represent the model within the substrate.

The two-dimensional substrate (Figure 2) employs only the x and y axes and is the simpler of the two substrates. The flow of information from the input layer to the output layer is along the y-axis. All input nodes are evenly distributed along $y = 0.0$; all hidden nodes are along $y = 0.5$. Finally output nodes are arranged along $y = 1.0$. In the output layer, the three left-most nodes represent the left leg effectors and the three right-most nodes represent the right leg effectors. Figure 2 shows the input layer with six inputs representing the six joint angles of the biped model; however, another configuration for the substrate is to eliminate all input nodes except nodes (0.167, 0.0) and (1.000, 0.0), which represents only whether the feet of the biped walker are touching the ground or not.

The three-dimensional substrate (Figure 3) employs the x, y and z axes.

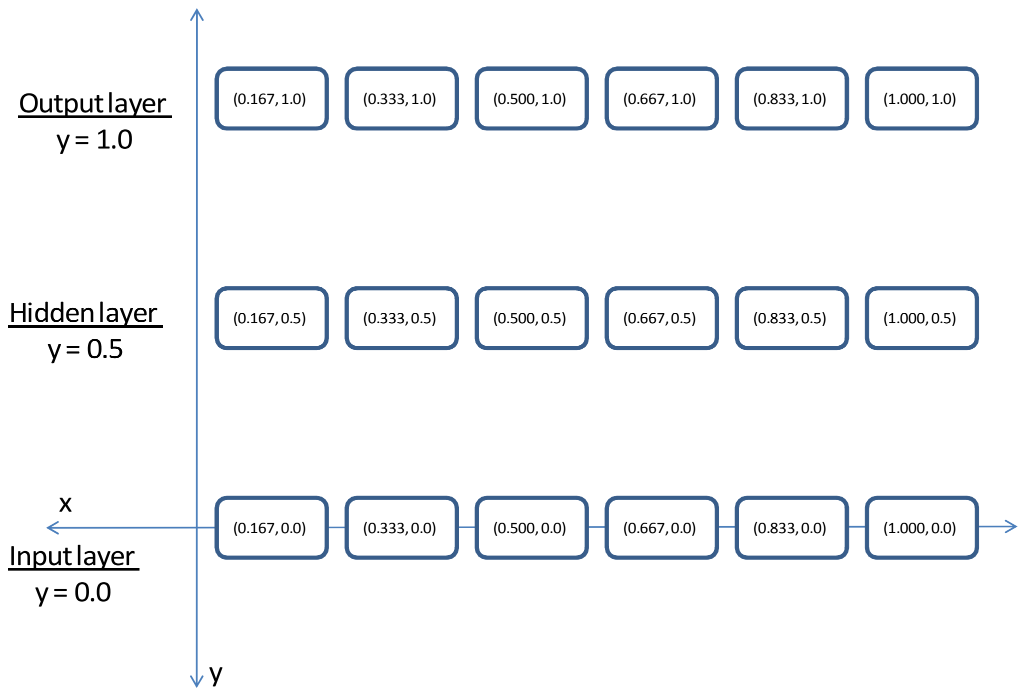

Figure 2: **Two-Dimensional Substrate.** A visualization of the two-dimensional substrate in the biped walker experiment is shown.

The flow of information from the input layer to the output layer is along the z-axis. All input nodes are distributed in the shape of the biped walker along $z = 0.0$; all hidden nodes are along $z = 0.5$ and all output nodes are along $z = 1.0$. In the output layer, the three left-most nodes represent the left leg output controls and the three right-most nodes represent the right leg output controls. Figure 3 shows the input layer with six inputs; however, as with the two-dimensional substrate, another configuration for the substrate is to eliminate all input nodes except nodes (-1.0, -1.0) and (1.0, -1.0) to represent only whether the feet of the biped walker are touching the ground or not.

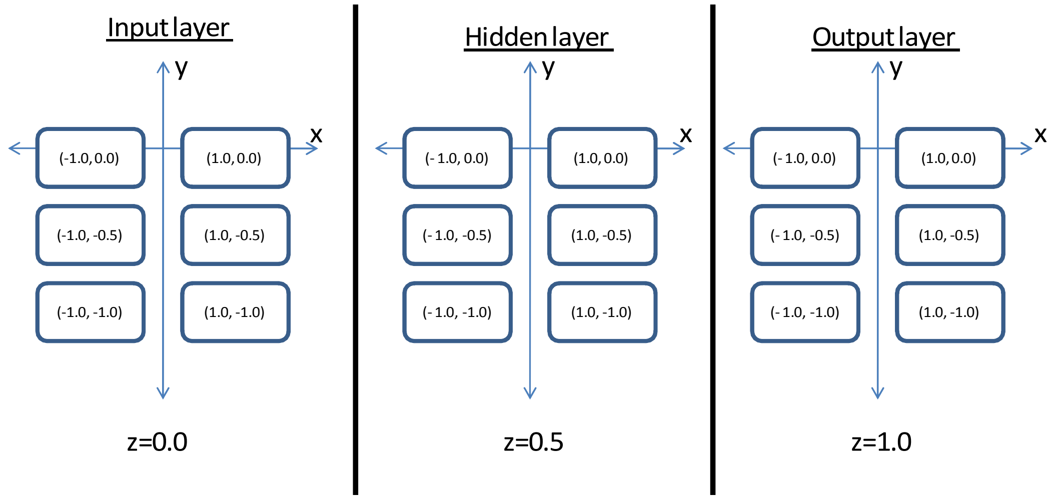

Figure 3: **Three-Dimensional Substrate.** The three-dimensional substrate for the biped walker experiment is depicted.

3.2 Functional Modularity

According to Wiegand et al. (2009), *functional modularity* is the localization of functionality according to structure. For example, if the human brain is examined while the human performs a specific action, a specific section of the brain, or *module*, is observed to correspond to that action (Bartels and Zeki 2005). It is widely believed that functional modularity plays an important role in the functional capacity of the human brain and in evolutionary computation. Thus it could be worthwhile to encourage functional modularity within the substrates of the experiment.

Functional modularity can be encouraged in HyperNEAT by modifying the *connection threshold* parameter; if a connection weight value between two nodes does not exceed the connection threshold, then the connection is not expressed. The connection threshold, T_{AB}, for two nodes A and B can be changed from a static value to:

$$T_{AB} = t_{min} + t_{mul} * E_{AB}, \tag{1}$$

where t_{min} is the minimum threshold value desired (for self-recurrent connections), t_{mul} is the threshold multiplier value, and E_{AB} is the Euclidean distance between nodes A and B. The effect is that if nodes A and B are close on the substrate, their connection is more likely to be expressed, whereas if nodes A and B are far apart, their connection is less likely to be expressed.

In the case of the two-dimensional substrate (Figure 2), only the x coordinate of the nodes is taken into account in the distance calculation because the only functional modularity to be discovered is along the y-axis (i.e. left leg controls vs. right leg controls). For the same reason, only the x and y coordinates of the nodes in the three-dimensional substrate (Figure 3) are taken into account in the distance calculation.

3.3 Novelty Search for Bipeds

Past experiments with NEAT have indicated that the regular fitness metrics are deceptive in the biped walker domain (Allen and Faloutsos 2009). While searching the biped walker behavior space seeking only the longest-walking gait may seem intuitive, it does not discover very effective gaits because regular fitness metrics become stuck on simple gaits that fall down quickly and are unable to elaborate upon them from there (Allen and Faloutsos 2009; Lehman and Stanley 2008). Thus, it may become necessary in the search to explore less-effective gaits in the hope that later elaborating upon them will result in a longer-walking gait than previously discovered.

One means of accomplishing this is searching for the most *novel* walking behaviors, that is, by awarding the highest fitness to the most unique walking behaviors in the current population amongst all walking behaviors ever discovered in the run (Lehman and Stanley 2008). This fitness metric is called novelty search and effectively evolved biped walkers with regular NEAT in informal experiments (unpublished; by Joel Lehman in the EPlex group). Thus instead of only looking at how far the biped walks, novelty search defines a *behavioral characterization* and instead explores the *behavior space* of the biped walker domain. Lehman and Stanley (2008) provides a more detailed description of the Novelty Search metric.

The same novelty search behavioral characterization used in the biped walker experiment with NEAT (by Joel Lehman in the EPlex group) is used in this experiment. In one second intervals throughout the simulation, the biped's center of mass is recorded according to the following equation:

$$x'_k = sign(x_k - x_0) * (x_k - x_0)^2 \tag{2}$$

$$y'_k = sign(y_k - y_0) * (y_k - y_0)^2, \tag{3}$$

where x_0 and y_0 are the biped's initial center of mass in the xy-plane and

x_k and y_k are the biped's center of mass recorded at simulation second k. If the biped falls at any point before the end of the simulation, the remaining x_k's and y_k's are set to the final x_k and y_k before it fell to the ground. Upon termination of the simulation, all of the recorded center of mass values (x'_i,y'_i), where 1$\leq$ i $\leq$ m, are concatenated into a list to form [(x'_1,y'_1),(x'_2,y'_2),...,(x'_m,y'_m)]. The biped's novelty value is then calculated by adding together all of the squared distances between the biped's behavioral characterization values (i.e. the center of mass records). It is important to note that two biped walkers that end up at the same location by means of a different path will receive two different novelty values. The two novelty values will be different because the behavioral characterization values are recorded throughout the simulation. These different novelty values then indicate that the two biped walkers are different from each other.

By defining biped walking behaviors in this way, novelty search does not necessarily search for the longest-walking behaviors. In fact, novelty search can indicate that a biped walking behavior that falls down quickly is just as good as a biped walking behavior that walks a long distance. Nevertheless, novelty search achieves the goal of continually exploring new, complex biped walking behaviors and can potentially discover a more stable walking behavior than normal fitness metrics by opening up the search space.

4 Experiments

This section describes the biped walker experiment and the attempted solutions.

4.1 Biped Walker Domain

The goal of the biped walker is to walk as far as possible from its starting point without falling down within the allotted amount of time. If at any point the biped walker falls down, the simulation for that biped walker controller is terminated and the maximum distance walked by the biped walker is recorded. To accomplish this sequence, the evolved ANN is given control of the six joints of the biped walker model and allowed to determine the angles as the simulation progresses, similar to the biped control model in Reil and Husbands (2002). To confirm the experiments previously done in the biped walker domain with NEAT (by Joel Lehman in the EPlex group), the NEAT experiment is attempted once again. Confirming that stable, long-walking gaits can be evolved by the NEAT experiment will validate that the novelty search approach is a good basis for implementing HyperNEAT in the biped walker domain.

Following the above control model, the substrates in this experiment are designed to express the geometry of the biped walker problem (Figure 4). All of the right leg joint nodes are placed to the right of the y-axis and all of the left leg joint nodes are placed to the left of the y-axis in the substrate,

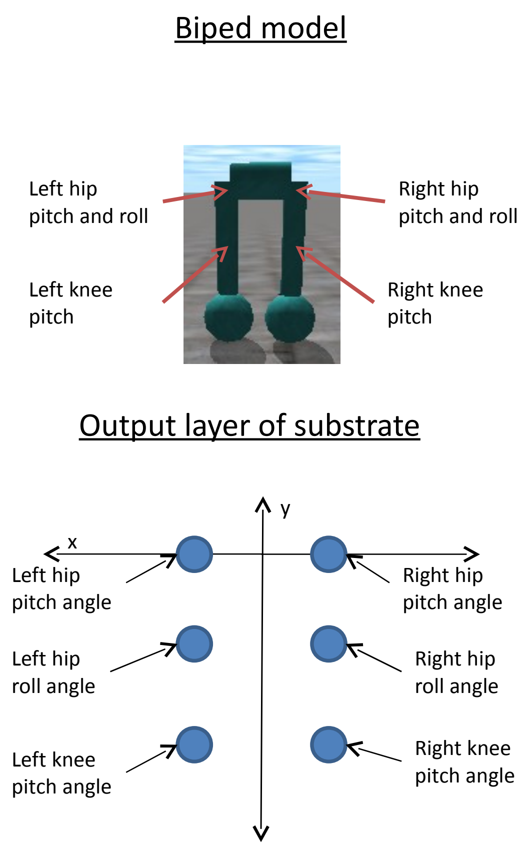

Figure 4: **Biped to Substrate Mapping.** The biped model and its joint locations are compared to the joint locations in the output layer of the substrate.

which maps directly to the biped model. All nodes are either on or below the x-axis to enforce that the only symmetry to be discovered in walking is between the two legs. For each leg, the two hip angle nodes are placed above the knee angle node, further encoding the geometry of the biped model in the substrate. Although it would be preferable to have both hip angle nodes for each leg in the same location, nodes are not able to occupy the same position in HyperNEAT substrates, so a compromise is chosen instead.

Beyond the design of the output layer of the substrate, many experimental design decisions are explored. Because the substrate link weights, biases and time constants are all generated by the CPPN, a range for each such parameter must be defined. The first option is to enforce common link weight, bias and time constant ranges that have been enforced in successful HyperNEAT experiments in the past (Stanley et al. 2009; Gauci and Stanley 2008, 2007). These common ranges may initially prove useful; however, better ranges for the biped walker problem may exist. To get a better idea of what these optimal ranges are, the longest-walking and most stable biped walkers from the NEAT-evolved biped walker experiment (by Joel Lehman in the EPlex group) are examined.

Because the biped walker problem has not been previously attempted with HyperNEAT, it is currently unknown what substrate configuration is best. As such, a wide array of substrate configurations can be explored. The geometry of the substrate can either be in two or three dimensions. Previous experiments in which HyperNEAT evolved quadruped walkers had success with three-dimensional substrates (Clune et al. 2009b), hinting that three-dimensional substrates may also work well for the biped walker problem. In this experiment, the two substrates described in Section 3.1 are tested.

Macro-level properties of the connectivity may also be significant. A *feed-forward* substrate is a substrate in which the links between nodes only point from the input layer to the hidden layer and from the hidden layer to the output layer; the flow of information is thus directly from the input layer to the output layer. In contrast, a *recurrent* substrate is a substrate in which links are allowed to connect backward, e.g. a node in the output layer is allowed to connect to a node in the hidden layer, or even to another node in the output layer. Both substrates can be configured as either feed-forward or recurrent. However, some level of recurrence is often required to create oscillatory behaviors, which are needed to walk.

Finally, the input layer of the substrate can be configured in a variety of ways depending on what information is to be given to the CTRNN controlling the biped walker. The six joint angles of the biped walker model can be given as inputs to the CTRNN, in which case the input layer can be designed exactly as the output layer in Figure 4. Alternatively, two foot touch sensors can be given as input; these foot touch sensors would input a value of 1.0 or 0.0, indicating whether their corresponding foot is currently touching the ground or not, respectively. These two substrate input configurations are tested in this thesis. In the previous biped walker experiment with NEAT, the two-input foot touch sensor input configuration performed most promisingly, perhaps indicating that this input configuration is best for the biped walker problem.

Following the precedent in the previous experiment with NEAT in the biped domain (by Joel Lehman in the EPlex group), fitness is compared to novelty in this case with *HyperNEAT* instead of NEAT. Fitness is how far the biped walker walked from its initial position, which is a common fitness measure amongst evolutionary biped walker experiments (Hein et al. 2008;

McHale and Husbands 2004; Reil and Husbands 2002; Van de Panne and Lamouret 1995). Distance traveled is intuitive to the biped walker problem as it proportionally rewards the biped walkers that move the furthest, which is an indication that the biped walker has a stable, oscillatory gait. This fitness measure will be compared to *novelty search*, described in Section 3.3.

Functional modularity is widely believed to play an important role in the functional capacity of the human brain and in evolutionary computation (Wiegand et al. 2009). Because it may apply well to the biped walker domain, in some experiments functional modularity is encouraged within the substrate, as described in Section 3.2, to see if it helps to evolve stable oscillatory gaits. By encouraging modularity within the substrate, it is hoped that the substrate can loosely break into two functional modules: a left leg module with the three left-most nodes of each layer, and a right leg module with the three right-most nodes of each layer. These modules would indicate that there is a strong correlation between the right-and-left leg inputs and the right-and-left leg outputs, respectively.

5 Results

From an early point in the experimentation, it became clear that the three-dimensional substrate, though logically designed, was not well-suited for the biped domain. No experiments were able to evolve a stable bipedal gait with the three-dimensional substrate, regardless of the parameters changed: input configuration, functional modularity, connectivity, fitness function, etc. In addition, champion walker genomes from the NEAT experiment mapped more clearly to the two-dimensional substrate when converted to a substrate. Due to these observations, the three-dimensional substrate was abandoned early in experimentation in favor of the two-dimensional substrate.

By studying the champion walkers from the NEAT experiments, a better parameter range for the HyperNEAT biped walker experiment can be determined. Table 1 lists the ranges that evolve the most stable biped walker gaits. These ranges were discovered by calculating the average minimum and maximum values of the link weights, biases and time constants of the biped walkers generated in the earlier NEAT-based biped walker experiment.

As expected, it became clear that a certain level of recurrence is required within an ANN to exhibit oscillatory behavior. Thus experiments with the feed-forward substrate were unable to successfully evolve oscillatory bipedal gaits, whereas substrates that allow recurrence were often able to exhibit at least oscillatory bipedal behavior. Of the substrate input configurations explored, only the two-input foot touch sensor configuration resulted in a

Parameter	Range
Link weight	[-5.0, 5.0]
Bias	[-4.0, 4.0]
Time constant	[0.1, 2.0]

Table 1: **HyperNEAT Parameter Ranges Used in Biped Walker Experiment.** The ranges are shown in which the link weights, biases and time constants were constrained when a substrate was generated by a CPPN in the biped walker experiment.

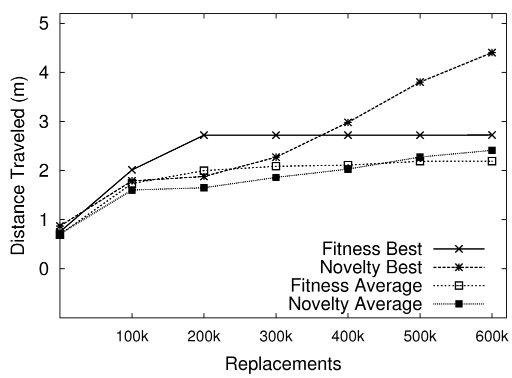

Figure 5: **Comparison of Fitness and Novelty Search in the Biped Walker Domain.** The average and farthest distance walked by biped walkers evolved with fitness and novelty search as a fitness metric are compared.

stable bipedal gait. This result confirms the findings in the previous biped walker experiment with NEAT (by Joel Lehman in the EPlex group).

Further confirming the findings of Joel Lehman in the previous biped walker experiment with NEAT, novelty search also evolves longer-walking bipeds in HyperNEAT. As demonstrated in Figure 5, novelty search is better able to explore the deceptive biped walker behavior space to discover longer-walking gaits than the fitness metric of distance walked. Though the fitness metric is able to discover biped walking behaviors which walk two to

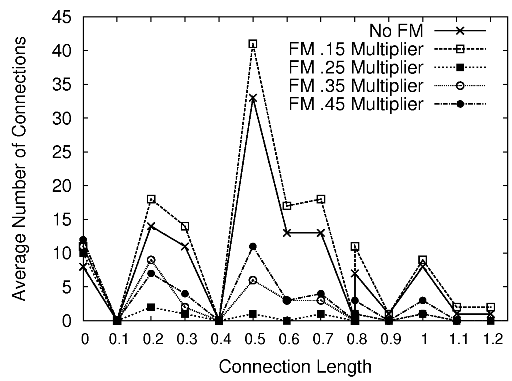

Figure 6: **Effect of Variable Connection Threshold on Connectivity.** The average number of connections of each length are plotted for each FM multiplier. *FM* indicates *Functional Modularity* and the number following it indicates the value of t_{mul} in the variable connection threshold equation. *No FM* indicates a static connection threshold of 0.3.

three meters at a faster rate than novelty search, the fitness metric quickly plateaus and is unable to improve upon the walking behaviors from there. In contrast, novelty search is capable of continually discovering new, farther-walking biped walking behaviors throughout the entire experiment.

To confirm that the functional modularity configuration is having the expected effect on the connectivity of the substrate, the average number of connections of each length are plotted in Figure 6. A t_{mul} value of 0.15 results in slightly more connections than the static threshold configuration overall, which is expected because the maximum value possible for E_{AB} is $\sqrt{2}$ in the two-dimensional substrate (i.e. the distance from the bottom left corner to the top right corner of the substrate). The low t_{mul} value creates a lower connection threshold, ranging from 0.15 (for self-recurrent connections) to 0.36, allowing connections to be more easily expressed. A higher t_{mul} value

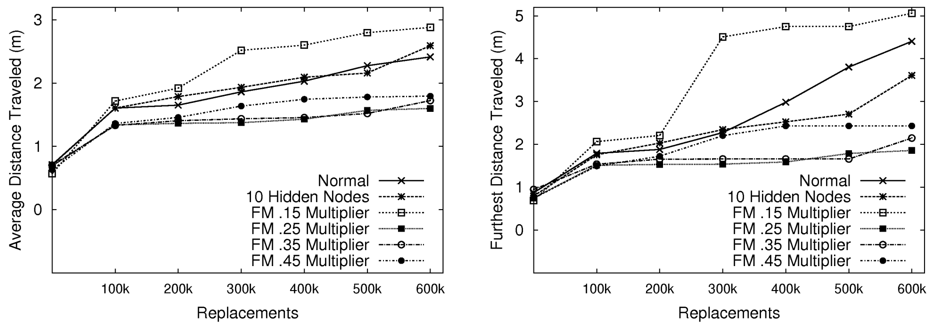

(a) Average Performance (b) Best Performance

Figure 7: **Performance Comparison of Substrate Configurations.** Various substrate configurations are compared based on the distance the evolved walker travels as learning progresses. *FM* indicates *functional modularity* and the number following it indicates the value of t_{mul} in the variable connection threshold equation. *10 Hidden Nodes* is a substrate configuration exactly the same as the *Normal* configuration, except with 10 hidden nodes in the hidden layer instead of only 6. Plot (a) compares the average distance traveled of the evolved biped walkers while (b) compares the distance traveled of the best walkers evolved for each configuration.

prohibits connectivity, although not in the manner expected. Interestingly, as the t_{mul} value increases from 0.15 to 0.25, the average number of connections drops significantly compared to the static threshold ($p < 0.01$), though as the t_{mul} value increases from 0.25 to 0.45, the average number of connections begins to rise again.

The above findings result in a two-dimensional, recurrent substrate with a two-input foot touch sensor input configuration; the link weights, biases and time constants are restricted to the ranges in Table 1. Novelty search is the fitness function. This substrate configuration is used as a basis for the biped performance comparison in Figure 7, labeled as the Normal configuration. Of the functional modularity experiments, only the experiment with the t_{mul}

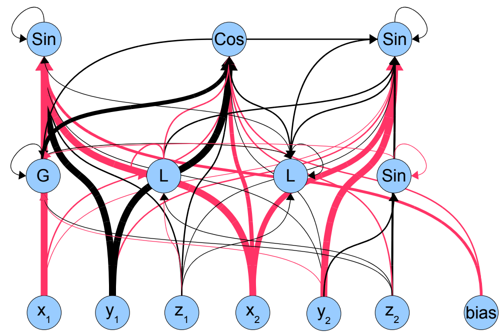

(a) CPPN

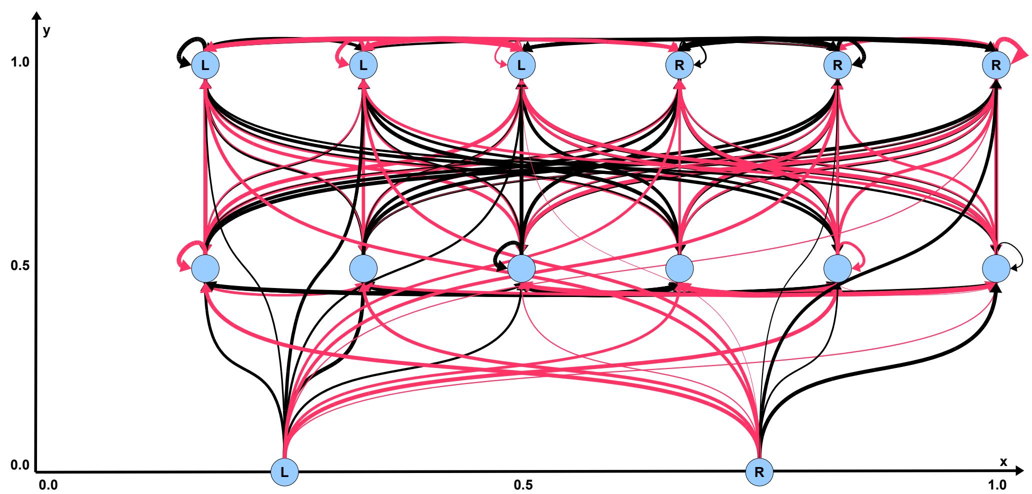

(b) Substrate

Figure 8: **Illustration of Champion Biped CPPN and Substrate.** The champion's CPPN is depicted in (a), where (x_1, y_1, z_1, x_2, y_2, z_2) are the inputs. The top three nodes are the outputs, from left to right: connection weight, node bias and node time constant. The activation function is the label inside the node, where Sin = Sine; Cos = Cosine; G = Gaussian; and L = Linear. The bias has a standard value of 1.0. The champion's substrate that is generated by the CPPN is shown in (b), where L and R represent nodes corresponding to Left and Right legs of the biped model, respectively. From left to right, the inputs are: Left-foot foot touch sensor, Right-foot foot touch sensor. The outputs, from left to right, are: Left knee pitch, Left hip roll, Left hip pitch, Right hip pitch, Right hip roll, Right knee pitch.

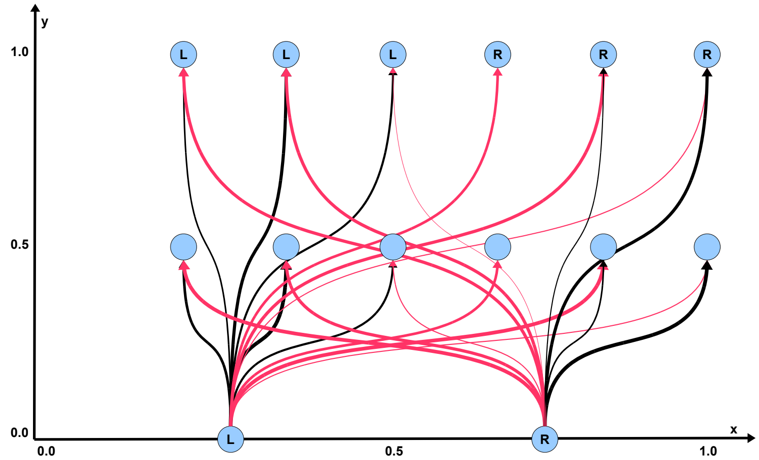

Figure 9: **Visualizing Symmetry in a Subset of the Champion Substrate's Connections.** A subset of the champion substrate's connections in Figure 8b is shown to more easily discern the symmetry within the substrate.

value of 0.15 did not hamper the performance of the Normal configuration. In fact, the functional modularity experiment with the t_{mul} value of 0.15 appears to perform slightly better than the Normal configuration both on average and in the best case, though the difference is not significant. There is seemingly no difference in the performance of the Normal configuration and the 10 Hidden Nodes configuration on average, though the Normal configuration performs somewhat better in the best case, though the difference is not significant. The main result is that in the best settings, HyperNEAT can evolve oscillatory gaits *and* balance for several meters, although the behavior is still brittle (i.e. it can fall easily).

Figure 8 illustrates the CPPN and substrate of the overall champion biped walker from all of the experiments, which is from the functional modularity experiment with a t_{mul} value of 0.15, whose fitness curve appears in Figure 7b. Figure 9 shows a subset of the full champion substrate to highlight the symmetry in its connectivity pattern. In particular, the left foot touch sensors pass positive activation to the left-side output nodes and negative information to right-side output nodes. The right-foot foot touch sensor connections exhibit the opposite pattern. Thus the substrate exploits the inherent symmetry of the biped, as was hoped for HyperNEAT.

In summary, the results of the HyperNEAT biped walker experiment indicate that the best substrate configuration is a two-dimensional, recurrent substrate with a two-input foot touch sensor input configuration and the link weights, biases and time constants restricted to the ranges in Table 1. Novelty search outperforms pure fitness as a fitness function in the HyperNEAT biped domain. The functional modularity equation affected substrate connectivity as expected by discouraging longer connections and encouraging shorter connections, though it had a negative effect on performance as the multiplier was increased. Changing the configuration of the hidden layer of the substrate had no notable effect on performance. In the best-performing controllers, HyperNEAT discovered symmetry between the left and right leg areas of the substrate.

6 Discussion

It is well-known that natural bipeds such as humans evolved through an indirect encoding, i.e. DNA (Gilbert 2000). In addition, bipedal walking has a regular pattern to it, which indirect encoding should be able to exploit. However, prior to this thesis, an indirect encoding has not been applied to the biped walker domain. In this thesis, I evolved a basic biped walking behavior with HyperNEAT. HyperNEAT generated oscillatory behavior, balanced a biped for over five meters, and evolved symmetries in the substrate. However, despite these achievements, the biped walkers evolved with HyperNEAT do not walk as far, on average, as the biped walkers that were evolved in the NEAT experiments, which employ a direct encoding.

This result is important because it exposes a weakness in indirect encoding that may not have been discovered prior to this thesis. It is difficult for an indirect encoding to precisely tune the connection weights of the ANN controller, which may be important in the biped walker domain. There is a certain degree of brittleness in the indirectly-encoded ANN controllers and they are highly susceptible to being rendered ineffective by small mutations. These observations are surprising because the same problem is not observed in nature: human beings experience some mutations every generation (Gilbert 2000) and are still capable of learning highly-tuned walking behaviors. Thus it may be that a key ingredient for the solution of the biped walker problem is missing.

Perhaps what is missing, then, is the ability to precisely tune the connection weights evolved by the indirect encoding *during the agent's lifetime*. Even evolved biped walkers that walk with a stable, oscillatory gait still exhibit slight imperfections in their gait that eventually build up and lead to losing balance and falling. If these imperfections were able to be mitigated *during the simulation* through adaptation, the controller would be able to dynamically fix them and maintain balance. Furthermore, the controller could potentially also adapt to new walking tasks during the simulation, such as walking up stairs for the first time or walking after a foot is damaged. There is evidence of such lifetime adaptation in nature: natural bipeds do not begin walking perfectly upon birth. Instead, they slowly adapt to their body and learn to walk efficiently as they grow into adulthood. Natural bipeds are also capable of adapting to changes in their body: if a natural biped experiences a physical change in its body, such as injuring its leg, it is able to adapt to this change and effectively walk with a limp. It is therefore likely that an adaptive mechanism needs to be added to the indirect encoding model to evolve more natural, longer-walking biped behaviors.

Another feature of the model that merits further investigation is the source of *feedback* to the biped controller during simulation. Instead of only receiving the two foot touch sensors as feedback, the biped controller could also benefit from a source of feedback indicating the model's current balance. Natural bipeds such as humans can perceive balance through the vestibular system; a simulated biped model could potentially benefit from a similar artificial vestibular system. Such a system could be simulated, e.g. by providing the biped controller with information about the model's current center of mass during the simulation. Two useful means of providing such information could be (1) to simply provide the biped model's current center of mass in the xy-plane or (2) to provide the distance from the biped model's current center of mass to a perfectly-balanced center of mass point. In either case, the biped walker controller would be constantly provided with relevant information about its current balance, allowing the controller to adjust accordingly.

The scope of this thesis is not limited to biped robots. This research is a preliminary step toward evolving fully-animated ragdolls that can potentially be deployed in video games and animated film. Instead of scripts controlling the models, ANN controllers may someday provide realistic, dynamic reactions to the model's surroundings and save animators the effort of creating an animation for every reaction. In a presentation in 2009, an Electronic Arts employee stated that the company wanted to move toward using ragdoll models in their Madden series (White 2009). Research such as that in this thesis will be necessary to realize this goal.

7 Conclusion

This thesis presented a new approach to evolving bipedal walking behaviors. With the combination of an indirect encoding, HyperNEAT, and novelty search to mitigate deception, basic biped walking behavior is evolved. Although indirect encoding would ideally evolve biped walking behavior as fluid and stable as that in nature, the results in this thesis suggest that such fluidity is beyond the current state of the art. Yet this gap suggests a possible future path through adding adaptive mechanisms to evolving stable, longer-walking biped behaviors.

References

Allen, B. and Faloutsos, P. (2009). Complex networks of simple neurons for bipedal locomotion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Bartels, A. and Zeki, S. (2005). Brain dynamics during natural viewing conditions–a new guide for mapping connectivity in vivo. *NeuroImage*, 24(2):339 – 349.

Bentley, P. J. and Kumar, S. (1999). Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-1999)*, pages 35–43.

Bongard, J. C. (2002). Evolving modular genetic regulatory networks. In *Proceedings of the 2002 Congress on Evolutionary Computation*.

Clune, J., Beckmann, B. E., Ofria, C., and Pennock, R. T. (2009a). Evolving coordinated quadruped gaits with the hyperneat generative encoding. In *CEC'09: Proceedings of the Eleventh conference on Congress on Evolutionary Computation*, pages 2764–2771, Piscataway, NJ, USA. IEEE Press.

Clune, J., Ofria, C., and Pennock, R. T. (2008). How a generative encoding fares as problem-regularity decreases. In *Proceedings of the 10th international conference on Parallel Problem Solving from Nature*, pages 358–367, Berlin, Heidelberg. Springer-Verlag.

Clune, J., Ofria, C., and Pennock, R. T. (2009b). The sensitivity of hyperneat to different geometric representations of a problem. In *GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 675–682, New York, NY, USA. ACM.

Gauci, J. and Stanley, K. O. (2007). Generating large-scale neural networks through discovering geometric regularities. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007)*, New York, NY. ACM Press.

Gauci, J. and Stanley, K. O. (2008). A case study on the critical role of geometric regularity in machine learning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*, Menlo Park, CA. AAAI Press.

Gauci, J. and Stanley, K. O. (2010). Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation*. To appear.

Gilbert, S. F. (2000). *Developmental Biology*. Sinauer Associates, Sunderland, MA, sixth edition.

Hein, D., Hild, M., and Berger, R. (2008). Evolution of biped walking using neural oscillators and physical simulation. *RoboCup 2007: Robot Soccer World Cup XI*, pages 433–440.

Hornby, G. S. and Pollack, J. B. (2001). The advantages of generative grammatical encodings for physical design. In *Proceedings of the 2002 Congress on Evolutionary Computation*.

Hornby, G. S. and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, 8(3).

Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*, Cambridge, MA. MIT Press.

McHale, G. and Husbands, P. (2004). GasNets and other evolvable neural networks applied to bipedal locomotion. In *From animals to animats 8: proceedings of the seventh [ie eighth] International Conference on Simulation of Adaptive Behavior*, page 163. The MIT Press.

Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Trans. Evolutionary Computation*, 6(2):159–168.

Reisinger, J., Stanley, K. O., and Miikkulainen, R. (2005). Towards an empirical measure of evolvability. In *Genetic and Evolutionary Computation Conference (GECCO2005) Workshop Program*, pages 257–264, Washington, D.C. ACM Press.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162.

Stanley, K. O., Bryant, B. D., and Miikkulainen, R. (2005). Evolving neural network agents in the NERO video game. In *Proceedings of the IEEE 2005 Symposium on Computational Intelligence and Games*.

Stanley, K. O., D'Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.

Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127.

Stanley, K. O. and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.

Van de Panne, M. and Lamouret, A. (1995). Guided optimization for balanced locomotion. In *Computer animation and simulation*, volume 95, pages 165–177. Citeseer.

White, D. (2009). Personal communication. April 15th, 2009.

Wiegand, R. P., Anil, G., Garibay, I. I., Garibay, O. O., and Wu, A. S. (2009). On the performance effects of unbiased module encapsulation. In *GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 1729–1736, New York, NY, USA. ACM.