Measuring NEAT Performance on a Variably Deceptive Bipedal Walking Task

Maddox Riley
CS81 Fall 2015
Professor Meeden

Abstract

The experiments described in this paper aim to create a bipedal walking task whose level of deception can be adjusted on a sliding scale. This task is considered deceptive because it uses as its objective measure of fitness the x-distance traversed by a simulated robot, measured from a point at a specified height on the robot's body. If this point is set to a substantial height above the robot's feet, simply falling flat, with the measurement point pivoting 90 degrees, increases the robot's objective fitness score substantially even as it fails the task. Falling is thus a local optimum. Traditionally, Novelty Search combined with NEAT would be thought to evolve a more effective Artificial Neural Network for such an agent than NEAT alone; Novelty Search tries to evolve substantially original solutions to a problem independently of the objective fitness measure, thus avoiding many of the local optima introduced by the type of deception described above. However, I wanted to compare the performance of NEAT alone with a simulated biped's fitness measurement point adjusted to various heights, creating varying levels of deception. The same experiment run using Novelty Search functions as a control. The goal is two-fold: to determine whether NEAT can produce similar success to Novelty Search on this comparatively difficult bipedal walking task; and if this is the case, to determine how NEAT's performance is affected by granular adjustments to the level of deception in its assigned task. Furthermore, I hope to demonstrate at which point a task might be considered deceptive enough that Novelty Search would offer substantial benefits compared to NEAT alone. Through my experiments, I was able to show increasing NEAT performance with decreasing deception by some metrics. However, these results were not entirely consistent and may not be entirely conclusive. The use of Novelty Search as a control failed to show satisfactory performance on this task. In fact, NEAT outperformed Novelty Search by a substantial margin for every experiment.

Introduction

The inspiration for this study comes from Joel Lehman and Kenneth O. Stanley's 2011 paper "Abandoning Objectives: Evolution through the Search for Novelty Alone". The paper concerns itself with a general discussion of the Novelty Search algorithm, which avoids falling into deceptive local optima by basing its evolutionary mechanism on behavioral novelty rather than on objective fitness (3, pg. 1). Several experiments in deceptive environments are discussed in which a combination of Novelty Search and NEAT significantly outperform fitness-based NEAT due to the latter's tendency to become trapped in local optima (3, pg. 2). One of these experiments involves a bipedal walking task. In this task, a simulated 3D robot controlled by an evolved ANN must walk as far as possible within a given amount of time (3, pg. 24). The agent's only inputs are two sensors indicating a foot in contact with the ground. The robot has two hip joints and two knee joints, but no ankle joints, with varying degrees of freedom in each joint. Fitness is calculated based on squared distance traversed by the robot based on the location of the robot's center of mass (3, pg. 25). My experimental setup varies in a number of ways, mostly for the sake of simplicity, reliability and controllability of the results. For example, my agent's only joints are at the hip. Furthermore, the agent's legs do not leave the ground, but "contact sensors" of the type in the original experiment are provided by measuring the leg angle of the agent (described in greater detail in the next section). There is no simulation of gravity per-se in my simulator; instead, the agent simply falls if it overextends its legs, so it needs to develop an oscillatory behavior if it is to make much progress. All in all, my setup and agent are much simpler than the one described in Lehman and Stanley's experiment. However, my assumption is that much of the deception in the bipedal task arises through the choice of measurement point on the agent's body. This is the main variable in my experiments, in contrast to the original.

Another practical study for the development of a bipedal robot points out that the majority of current bipedal robots operate based on manually determined commands or pre-defined walking programs that are not particularly adaptive (2, pg. 2). This is very far-removed from the nature of Lehman and Stanley's experiment, and of the experiments I aim to perform. This distinction underscores the difficulty of trying to evolve an ANN that can control a bipedal agent with no biases, no pre-defined behavior, and only ground sensors on each foot (and, in my case, a countdown timer) as inputs to the evolutionary algorithm.

An article about a research team trying to model the mechanics of human bipedalism in a complex, multi-jointed robot highlights some of the specific difficulties in the bipedal walking task. For example, the joints often move at highly variable velocities; this is especially true in the example given in the article of knee joint movement in uphill walking or walking up a flight of stairs, though it applies more generally as well (1). In contrast, my simulated environment ignores much of this complexity, largely out of simple necessity. Another observation is that humans naturally tilt forward, potentially coming close to falling, when walking. This is another reason why the local optimum of pitching over is so easy to fall into (1). Though my simulator does not directly replicate this type of leaning behavior, the necessity of extending a leg, and thus moving closer to the dangerous angle at which falling over becomes a possibility, should mirror the deception of needing to lean over, risking a fall, in order to move forward.

Hypothesis

For my experiments, I predict that reducing the amount of deception in the bipedal walking task by lowering the point of distance measurement on the agent's body will result in a substantial improvement in the performance of NEAT on this task. I would still not be surprised to see Novelty Search outperform NEAT across the board, but I would expect Novelty Search's advantage to become substantial only when the task is made more deceptive by raising the measurement point.

Experimental Setup

All experiments were run using simulated bipedal agents. For the simulations, I coded a modified version of the Python and Zelle Graphics-based robot simulator used for a number of in-class Adaptive Robotics labs. The environment is always represented as a rectangular space of 1400 pixels in width by 400 pixels in height. An agent is represented as a set of five points: a head, a waist, a left foot, a right foot, and a point from which the current x position of the robot is measured. For my experiments, agents were defined to be 200 pixels in height and were placed directly in the center of the environment. The height of measurement point on the robot's body may be defined on a floating point scale from 0.0 (on the ground of the virtual environment) to 1.0 (at the very top of the agent, i.e. 200 pixels from the ground). Each foot is allowed a range of motion from -60 to 60 degrees about the agent's waist; the rest of the agent's body does not pivot. There are two motor outputs to the agent. The first is a floating point value between -1.0 and 1.0, where -1.0 - 0.0 inclusive specifies that the left foot should be rotated, and a value of 0.1-1.0 inclusive specifies a rotation of the right foot. The second is a floating point value between -1.0 and 1.0 indicating the degree of rotation about the waist, where value*60 = degrees of rotation. When one foot is rotated, the other foot is rotated an equal amount in the opposite direction. The agent moves to the left if a left foot rotation and a negative rotation amount are specified, and right if a right foot rotation and a positive rotation amount are specified. The movement is a translation of the whole agent in the x direction of 15 pixels times the normalized rotation amount from -1.0 to 1.0. In order to move both feet back towards the center of the agent, the controller must specify either a negative right foot rotation or a positive left foot rotation; this is necessary to develop a successful oscillatory walking behavior, as the agent will "fall" and its run will terminate if its feet are pivoted at an angle greater to or equal than 58 degrees. The agent could be said to move forward with a shuffling, scissor-like movement. When the agent falls, its entire body does not pivot, but the fitness measurement point pivots 89 degrees in the direction of travel, adding the element of deception to the walking task. A rotation of 89 degrees rather than 90 degrees is used so that the measurement point is still visible in the simulator window after an agent has fallen. Also for visualization purposes, lines are drawn between the head an waist, waist and left foot, and waist and right foot of the agent, representing a torso and legs. For clarity, the left leg is colored green and the right leg is colored red.

There are three sensory inputs to both NEAT and Novelty Search: the number of time steps remaining, initialized at 1000 in all experiments, and two sensors that read 1.0 if a foot is within 5 degrees of its "starting" position in-line with the rest of the agent's body and 0.0 if it is not. This is analogous to the sensors in Lehman and Stanley's experiments indicating whether or not a robot's foot is in contact with the ground.

Objective fitness is specified as a value between 0.0 and 1.0, a proportion of an agent's maximum possible distance away from the center of the world, that is, 700 pixels in either direction given a world of width 1400. Given these values, objective fitness as a function of the measurement point's x value is given by the equation:

abs(measure point X - 700) / 700.

Strict upper and lower bounds of +/- 700 are set on the measured distance away from the world center to account for cases where the measurement point may pivot outside the bounds of the agent's world.

For the Novelty Search experiments, the behavior space of an agent is defined as a series of x regions visited by either foot of the agent. Each region is 10 pixels in length, resulting in 140 total regions with values ranging from 0 for the leftmost region to 140 for the rightmost region. Behaviors are thus lists of 140 point of length 1, the list being padded with 0s in place of unvisited points. This behavior setup leads to a maximum length of 140*140 (140 points of length 1 with a maximum cartesian distance of 140 between any two points).

For all experiments, runs of 15 generations were used. For the Novelty Search experiments, the primary parameters were k=15 (k-nearest neighbors), a limit of 100, a threshold of 0.25, a point length of 1, and a maximum distance of 140*140. For further NEAT and Novelty Search parameters, see Appendix 1 at the end of the paper.

For NEAT, 7 total experiments of 10 trials each were run with fitness measurement points positioned at 0%, 25%, 33%, 50%, 67%, 75%, and 100% of the robot's height. For Novelty Search, the measurement point was set to 50% of the robot's height to approximate the conditions in Lehman and Stanley's paper.

Figure 1 (below): An example of an agent of height 200. The point in the center of the agent's body is the fitness measurement point, at 50% of the agent's height. The agent begins at its starting position (left), then extends its right leg 30 degrees, causing it to move forward slightly (center). It then extends its right leg another 30 degrees, again causing it to move forward, but because it has overextended, its measure point pivots 89 degrees and this run ends. The pivoting of the measure point increases the measure point's x value, making this situation a local optimum. To successfully traverse its world, the agent needs an oscillatory, shuffling, "scissor-like" walking movement

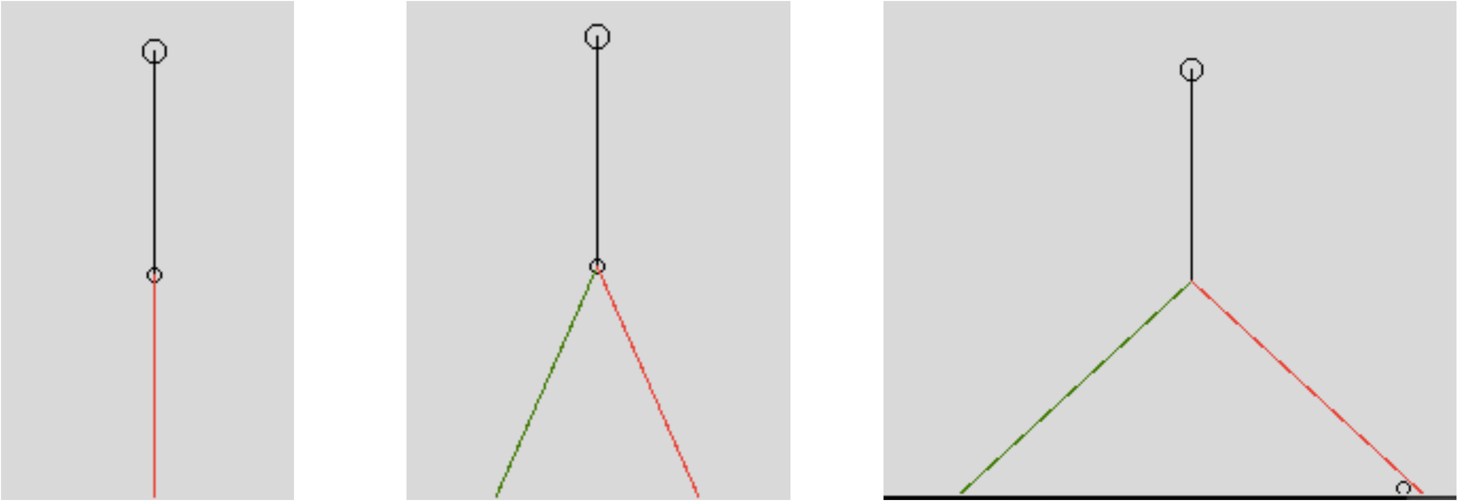

Results

The results of the experiments for various fitness measurement heights with NEAT as the evolutionary algorithm are summarized in Table 1 below:

Table 1
NEAT Experiment: Averages over 10 Runs per Measure point

Measure Point	0.0	0.25	0.33	0.5	0.67	0.75	1.0
Best Fitness	0.727	0.716	0.562	0.605	0.601	0.706	0.581
Generations to achieve best fitness (0-14)	6	5	5	7	5	5	7

In many cases, the average best fitness scores produced over ten runs are higher when the measurement point is set lower on the agent's body, thereby theoretically reducing the amount of deception in the walking task. This might seem to support my hypothesis that NEAT's performance on a bipedal walking task can be improved by lowering the amount of deception in the task. However, this is not always the case. Most notably, the average score for runs with the point set to 75% of the agent's height is notably higher than the averages scores for runs with the point set from 33-67% of the agent's height. Furthermore, the average fitness score for runs using a height of 0% is very similar to that of runs using a height of 25%, and there is a slight increase in the score from 33% to 50%. It is also difficult to notice a consistent trend in the number of generations needed to reach maximum best and average fitnesses, although peak average fitness scores often increase somewhat with theoretically increased deception, seemingly contradicting my hypothesis. Qualitatively, every run of NEAT developed at least some agents that were able to discover the necessary oscillatory shuffling behavior to completely traverse half of the world, thus achieving perfect fitness scores.

The results from the experiments using Novelty Search are summarized in Table 2 below:

Table 2
Novelty Experiment: Averages over 20 Runs total

Best Objective Fitness	0.191
Generations to achieve best objective fitness	5
Best Average Fitness (novelty)	0.101
Generations to achieve best average novelty	0

It is difficult to draw firm conclusions from these results. Novelty Search performs quite poorly on this bipedal walking task, rarely generating behaviors that cover more than 20% of the objective space. Qualitatively, the agents controlled by Novelty Search seem to move only a step or two before falling flat, consistently producing excessively large forward leg movements without counterbalancing reverse movements; oscillatory behavior is not discovered.

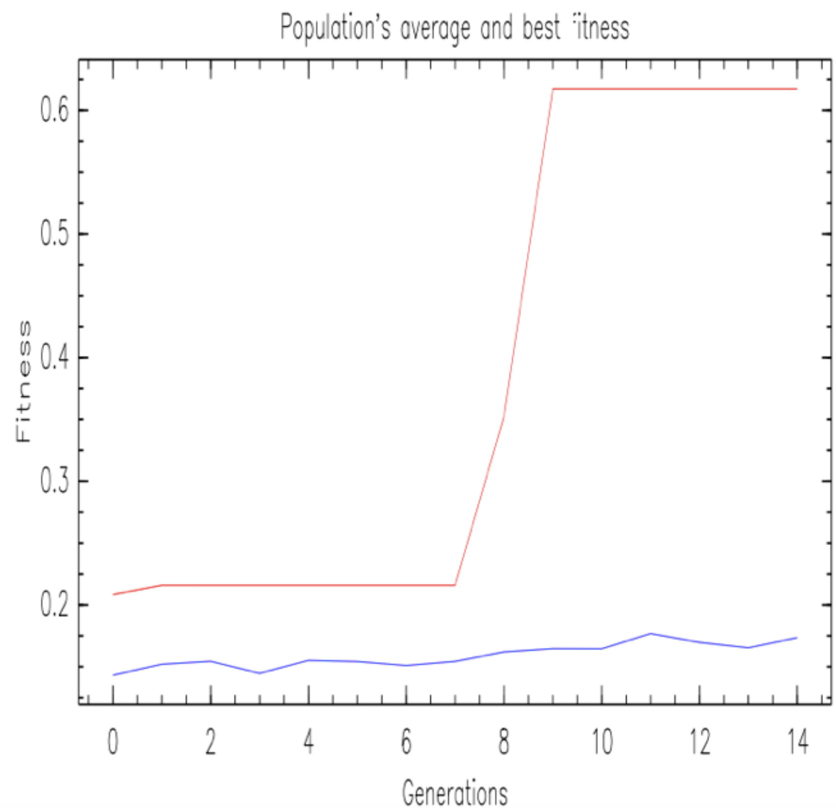

Figure 3: Average and best fitness scores from a fairly representative NEAT run using a measure point of 50%. The best score for a particular run will often converge to a particular value very quickly and/or suddenly.

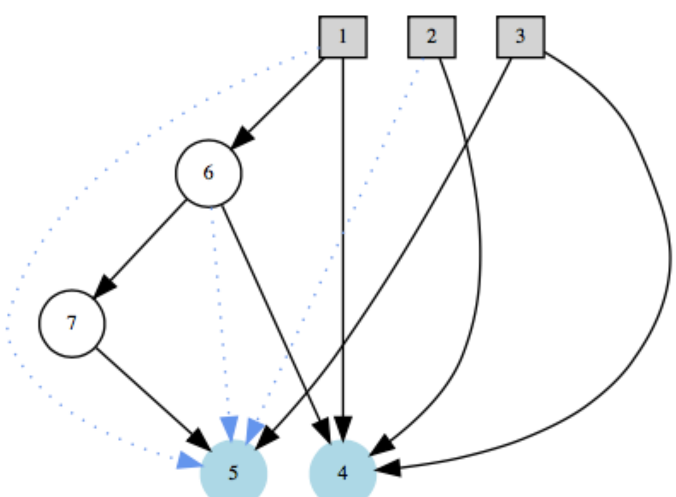

Figure 4: The phenotype of the best chromosome from generation 14 of a different run of NEAT, again with the measure point at 50%.

Conclusions/Discussion

Overall, I was not able to substantiate my hypothesis. While the results for the NEAT experiments may demonstrate some support to the prediction that varying the agent's point of fitness measurement will substantially affect NEAT performance, this did not hold true for every experiment, nor was there a consistent trend in most of the data points. In fact, NEAT converged to a correct solution very quickly in some cases. Perhaps more surprising was the extremely poor performance of Novelty Search on the same task. The higher measurement point in itself may partially explain the higher scores for some of the experiments run with the point above the robot's center. Although this change increases the deceptiveness of the task, it also leads to a larger x-distance increase should the agent fall. However, I think the main takeaway from my results is that the task of creating a bipedal walker in the real world, in three dimensions, with multiple joints, and with the effects of gravity, friction, and other factors, is far more complex that what I am reasonably equipped to recreate at this point. There are likely other sources of deception in the task than simply the choice of distance measure. This may explain why modifying this parameter did not yield quite the expected performance differences over the course of the NEAT experiments. In the case of Novelty Search, the algorithm may not have been able to perform to its full potential given what may not have been a highly deceptive task. Perhaps this is a demonstration that, for simpler tasks, using behavioral novelty rather than objective fitness can be detrimental, though this is only speculation on my part.

Appendix 1: Additional Parameters

NEAT Parameters	Novelty Search Parameters

[phenotype]
input_nodes = 3
output_nodes = 2
max_weight = 30
min_weight = -30
feedforward = 1
nn_activation = tanh
hidden_nodes = 0
weight_stdev = 0.9
 	
[genetic]
pop_size = 100
max_fitness_threshold = 1
 	
# Human reasoning 	
prob_addconn = 0.1
prob_addnode = 0.05
prob_mutatebias = 0.2
bias_mutation_power = 0.5
prob_mutate_weight = 0.9
weight_mutation_power = 1.5
prob_togglelink = 0.01
elitism = 1
 	
[genotype compatibility]
compatibility_threshold = 3.0
compatibility_change = 0.0
excess_coeficient = 1.0
disjoint_coeficient = 1.0
weight_coeficient = 0.4
 	
[species]
species_size = 16
survival_threshold = 0.2
old_threshold = 30
youth_threshold = 10
old_penalty = 0.2
youth_boost = 1.2
max_stagnation = 15
	[phenotype]
	input_nodes = 3
	output_nodes = 2
	max_weight = 30
	min_weight = -30
	feedforward = 1
	nn_activation = tanh
	hidden_nodes = 0
	weight_stdev = 0.9
	
	[genetic]
	pop_size = 10
	max_fitness_threshold = 1.1
	
	# Human reasoning
	prob_addconn = 0.1
	prob_addnode = 0.05
	prob_mutatebias = 0.2
	bias_mutation_power = 0.5
	prob_mutate_weight = 0.9
	weight_mutation_power = 1.5
	prob_togglelink = 0.01
	elitism = 0
	
	[genotype compatibility]
	compatibility_threshold = 3.0
	compatibility_change = 0.0
	excess_coeficient = 1.0
	disjoint_coeficient = 1.0
	weight_coeficient = 0.4
	
	[species]
	species_size = 10
	survival_threshold = 0.2
	old_threshold = 30
	youth_threshold = 10
	old_penalty = 0.2
	youth_boost = 1.2
	max_stagnation = 15

Bibliography

1. Hoopes, Heidi. “AMBER 2 bipedal robot mimics humans for a more natural gait”. *Gizmag*. 3 November, 2013. Online <http://www.gizmag.com/bipedal-robotic-locomotion-amber-2/29603/>

2. Kim, Jung-Kup et. al. “Experimental Realization of Dynamic Walking for a Human-Riding Biped Robot, HUBO FX-1”. *Advanced Robotics (RSJ)*. Accessed online. <http://www.cs.cmu.edu/~cga/legs/Paper_2.pdf>

3. Lehman, Joel, and Kenneth O. Stanley. “Abandoning Objectives: Evolution through the Search for Novelty Alone”. *Evolutionary Computation* 19.2 (2011): 189-223.