Observations on the Evolution of Neural Networks for the Control of a Simulated Quadruped Robot

Keith Wiley

University of New Mexico, Albuquerque, NM 87131

Abstract

This study used a genetic algorithm to evolve neural network controllers for the walking behavior of a quadrupedal robot. A 3D physics simulator (Breve by Jon Klein, using the ODE for physics simulation) is used to simulate the action of a quadruped creature in a 3D world. A variety of neural network architectures are experimented with, using a genetic algorithm to evolve walking behaviors that are rewarded with a fitness proportional to the horizontal distance traversed in a particular direction in a set time-period. Several different experiments test the effects of varying the neural network architecture, the friction of the ground, the pitch of the ground, and the strength of the creature's muscles. This paper demonstrates that the evolutionary system developed in this research is robust to a wide range of these variables, producing impressive walkers across a broad spectrum of physical and neurological circumstances.

Introduction

The nascent robotics industry is poised to expand rapidly in the early decades of the twenty-first century, with a wide variety of recognized applications that are ready for the introduction of physical robots capable of performing useful tasks in the real world. Many challenges still face the field of real-world robotics however. One such challenge is the successful production of agile legged robots. Only legged robots have any chance of functioning in arbitrary environments such as those presented by most situations in which people interact every day, and in which robots will have to function if they are to be useful.

This paper conducts a number of experiments on the generation of neural network locomotion controllers for a quadruped robot using a genetic algorithm as the means of behavioral development and a 3D physics simulator as the environment. While it is widely agreed that final results in this field will require implementation and demonstration on an actual physical robot, there is still a lot to be learned at the simulation stage of experimentation. Although it may be more difficult to produce final behavioral algorithms ready for actual physical robots using simulation — due to the imprecise representation of complex environments in a simulation — it is in fact easier to flesh out general approaches to questions concerning which methods might most successfully produce final results. Simulation allows numerous different approaches to be tried quickly, and with little cost, such that time spent on a physical platform at a later stage will be minimized, increasing efficiency and decreasing cost.

Breve

The experiments described in this paper were implemented using the Breve package for Mac OS X and Linux, created by Jon Klein (Klein, 2003). Breve was specifically designed as a tool for artificial life research, consisting of the coupling of multiple components, which when brought together, provide the necessary means by which to quickly and efficiently construct experiments in the physical modeling and rendering of artificial creatures. The salient components of Breve are:

- A 3D renderer for visualization of an animated 3D world (implemented using OpenGL)
- A physics engine (the open source ODE) capable of modeling the kinematics and dynamics of articulated creatures, in addition to modeling the effects of physical forces on such creatures and the effects of collision interaction between pairs of creatures
- An interpreted scripting language called Steve that enables fast implementation of and changes to a simulation from within Breve without the need to recompile code.

Since Breve is a physics simulator, it is possible to speculate on its running time in relation to "real" time. Actually, even with an accurate physics model such as Breve, the only physics parameter which truly corresponds directly to real time is the gravitational constant. Nevertheless, this single parameter brings the entire physics model into a direct relationship with real time. On a 1.6 GHz G5 PowerMac (circa 2003) Breve runs the model used for these experiments at approximately seven times real speed. For this reason, experimentation must be sparse. It is impossible to run thousands, much less tens of thousands, of trials on a given set of parameters with the intent of reducing noise and statistical error.

The Body Plan

The precise design of the body plan on which controllers would be developed was somewhat arbitrary. A quadruped was desired, primarily because it is the platform that receives the least attention in the research community, hexapods being popular for their inherent stability and bipeds being popular for their inherent human-like qualities. The basic shape and articulation was chosen to vaguely resemble the commercially available Sony AIBO robot, although with four fewer degrees of freedom (figure 1). Each leg has a hip and a knee, both of which rotate around the Z-axis where the robot faces along the X-axis and the Y-axis faces up (so the legs swing backwards and forwards along the side of the body). We chose to append a "head" was to the body, originally for aesthetic purposes. However, the presence of a head turned out to introduce interesting problems to the robot. It turns out that the controllers for the front legs evolve to be stronger in order to hold up the heavy head, an unexpected albeit somewhat obvious result. The head also introduces front/rear balance problems when the robot is placed on a pitched surface.

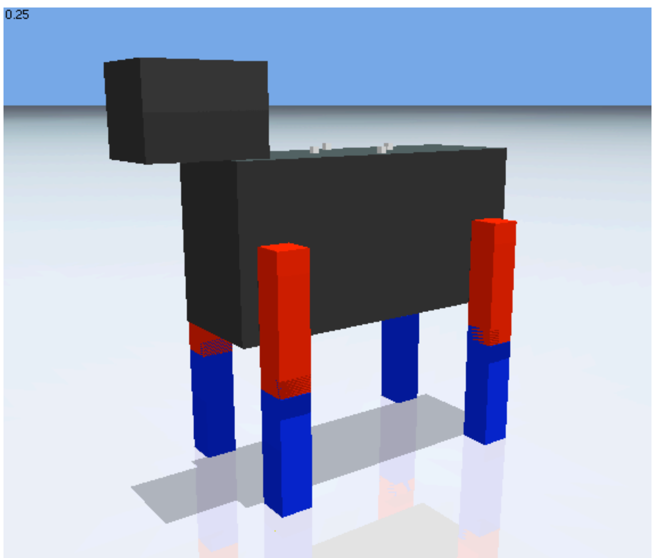

Figure 1: A BREVE screenshot showing the body plan we designed for this research, which vaguely resembles the Sony AIBO robot.

Neural Network Architecture

There are many ways in which a walking robot might be controlled. Some examples include subsumption architecture based algorithms, sets of sinusoidal oscillators, simple periodic signals, genetic programming trees, and neural networks. Some of these methods can be brought together as well. For example, the layers of a subsumption architecture based approach might be implemented as neural networks, or a set of sinusoidal oscillators might use sin parameters that are modified on the fly by genetic programming trees.

For this research we focused solely on neural network control, but with an externally imposed square-wave oscillating input signal to the neural networks. We cannot describe in simple terms the full structure of the neural networks here because part of the experimentation consisted of varying this structure, as is described later in this paper. Consistencies throughout the experiments include the use of sigmoid function nodes, the use of a bias node in the input layer, and the application of a square wave oscillating signal to a specified input node. The period of the oscillator was genetically determined in most experiments.

Additionally, we discovered early in this research that using multiple parallel neural networks provides a significant improvement to the performance of the walker and ultimately the genetic algorithm. Originally, a single neural network produced a set of outputs which drove the target velocities of the joints. The new approach produces four phenotypic neural networks from the genome, each network driving only a single leg. The genotype contains a description of only two neural networks, one for the front legs and one for the rear legs. It was assumed outright that bilateral symmetry in the controller should be an obvious quality of a bilateral body plan, and therefore a genotype with two neural networks is sufficient to encode the phenotype for four leg controllers in a quadruped.

We cannot state precisely what the inputs were in all the experiments, because, as with the structure (such as the number of layers, the number of hidden nodes, and the placement of bias nodes), the set of input nodes also varied with the experiments. However, we can list the maximum set of inputs that we defined:

- front/rear tilt sensor
- left/right tilt sensor
- body height above ground
- joint position for four hips
- joint position for four knees
- joint velocity for four hips
- joint velocity for four knees
- one oscillator signal
- limb length
- floor friction
- joint strength

The last three signals were only used in a suite of experiments in which the corresponding parameters could vary. In most experiments these inputs were not included at all. Likewise, the ground height did not turn out to be very useful and the tilt sensors were used in some, but not all, of the experiments. The precise combination of joint positions and velocities that were available varied somewhat in the earlier experiments until we settled upon a minimal yet completely functional set, as described below.

Genetic Algorithm

Most experiments were conducted with only minor variation to the genetic algorithm. In most cases, the population consisted of thirty controllers, each of which was tested on the simulated body for twenty simulated seconds. In later experiments, the trial period was reduced to ten simulated seconds. Fitness in almost all experiments was calculated as the square of the X component of the final displacement of the body from its starting position (only forward motion counted, backward motion resulted in a score of zero). Squaring the distance provides a disproportionate reward for controllers that traverse a greater distance. Rewarding only the X component of the XZ final position emphasizes efficient motion in a straight line. Reproduction was asexual, with parents chosen proportionate to fitness. Due to the small population, an elitism of ten percent was used to prevent the top three controllers from being lost should the dice “fall badly” on any given generation. Depending on the experiment being run at any given time, runs were terminated after some number of generations. In most cases this termination was triggered manually after the apparent plateau in evolutionary performance. In some cases, a large batch of experiments would all be automatically terminated after a preset generation in order to normalize any comparison made on the final results.

Evolution of Standing

At the outset of this research, it was unclear if the initial design would work at all. A neural network architecture had been chosen, but its design was fairly arbitrary, and it remained to be seen whether the system would permit reasonable evolution at all. We therefore decided to start with a presumably simpler task than walking, such as standing. At this point, the idea of producing multiple bilaterally symmetrical phenotypes had not occurred, so the phenotype consisted of a single neural network.

As it turned out, even the simple task of standing up could not confidently be evolved without some mechanism for built in symmetry. Multiple methods of symmetry were attempted. One such method consisted of a single neural network with two outputs, such that one output was explicitly fed to both front legs simultaneously and the other output was similarly fed to the rear legs. Although this approach worked, it was disappointing to have such an artificial mechanism in place.

The next method we attempted was altered very little for the rest of experimentation throughout this research. It consisted of describing two neural networks in the genome, one for a front leg and one for a rear leg, and producing two phenotypic neural networks from each, such that in the phenotype each leg gets its own signal processing network and output signal. At first glance this may not seem any less artificially imposed than the original method, but this method is quite different. Since each leg gets a unique signal, this system is at least ready for an increase in complexity which will allow left/right differentiation of leg behavior, where the original method would never have allowed such an extrapolation. While this is unnecessary for standing, it is vital for most locomotion gaits. Nevertheless, if two neural networks corresponding to the front left and front right legs are identical, then there is no hope of the two legs ever behaving differently, unless the inputs are capable of unique patterns as well. This point is extended further in the next section.

It is interesting to note that even an action as simple as standing can be performed in multiple ways and our evolutionary system is quite capable of finding multiple ways of accomplishing the task (figure 2).

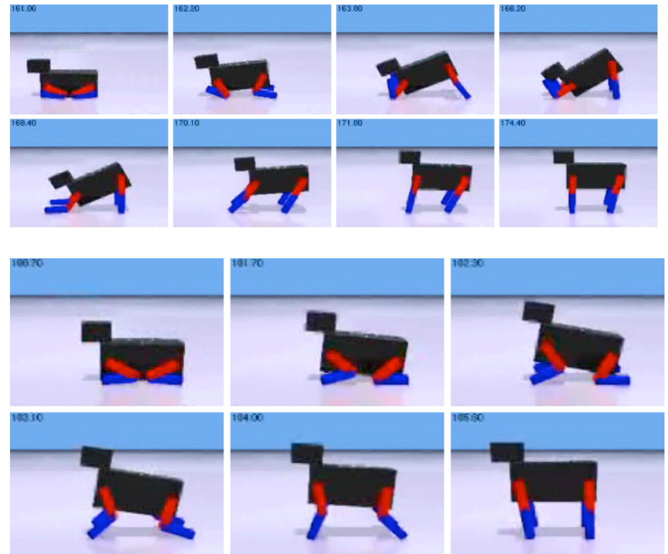

Figure 2: Two different evolved methods of standing

Evolution of Walking

In the previous section we pointed out that it would be necessary for the left and right legs to receive different inputs if there would be any hope that the left and right legs would receive different outputs. On perfectly flat ground, there is no conceivable source of variation in the input. This is not necessarily true if the ground is rough, in which case the inputs might reflect the uneven ground. However, this does not very easily allow for the cyclic leg patterns that are observed in all legged animals, to say nothing of the locomotion of nonlegged animals which predates legged locomotion by millions of years and also employs cyclic and oscillating muscle patterns. Clearly, the application of direct cyclic signals is justified by nature’s universal adoption of the same strategy.

It has been established that virtually all animals use periodic neural circuits called central pattern generators (CPGs) to coordinate the motion of their muscles during highly repetitive actions (Golubitsky et al, 1999; Ijspeert and Kodjabachian, 1999). In particular, CPGs are known to exist in the walking behaviors of legged animals (Lewis, 2002). While there has been extensive research into the construction and function of neural networks that act as CPGs (Chiel et al., 1999; Ijspeert and Kodjabachian, 1999), we did not want that to be the focus of our research, so we abstracted away the entire notion of a CPG, and reduced it to a single set of artificially imposed oscillating signals that were directly imposed on a single input node of each of the four phenotypic neural networks.

The oscillating signal consists of a square wave composed in equal parts of a series of +1s followed by a series of -1s. The square wave and positive/negative characteristics of this signal were not chosen arbitrarily. A square wave was chosen in favor of a — perhaps more intuitive — sin wave because it is known that the leg swinging behavior of some animals, insects in particular, is not composed of a sinusoidal set of neurological signals, but rather is composed of distinct forward and reverse swinging phases, which are triggered and set into continuous motion at the extreme positions of the stepping cycle. These extreme positions are commonly referred to as the Posterior Extreme Position (PEP) and the Anterior Extreme Position (AEP) (Barnes, 1998). For this reason, the oscillating signal we impose is broken into two distinct phases instead of a smooth and constantly changing function, although for reasons discussed below it is somewhat inappropriate to label either the positive phase or the negative phase as the "forward" or "reverse" phase of the leg stepping cycle. The choice of oscillating between +1s and -1s as opposed to some other likely signal such as 1s and 0s was made primarily to aid the reversal of direction of a leg by feeding signals of opposing sign into the neural network's input layer. This eases the task of evolving muscle motions which must completely reverse direction in order to produce a successful stepping motion. We did not explore the ability of our system to discover efficient oscillating leg motions with a different kind of oscillating input signal, although this is an interesting question, perhaps for future research. Clearly, there is no biological analogy to having a neural output signal that is negative, since biological neural signals are traditionally measured as the firing rate of a given neuron, which by definition must be positive, or in the most extreme case might be equal to zero. Nevertheless, in artificial neural network research, the usage of negative neural signals is quite common, generally for reasons similar to those stated above.

The oscillating signal has three parameters governed by the genotype:

- period
- left/right phase offset
- front/rear phase offset

The front left leg's network receives the bare oscillating signal as input, without any alteration from the simulation's internal clock. The front right leg receives the same signal offset by the left/right phase offset while the front rear signal receives the offset of the front/rear phase offset, and finally the rear right leg receives the signal offset by the sum of both phase offsets.

Initially, we were unsure how much input information would be necessary for the set of neural networks to walk successfully, so our predilection was towards giving the network as much information as we imagined might be useful. For this reason, each of the four neural networks in our first design had thirteen inputs, although bear in mind that evolution only had to coordinate twenty-six such inputs, not fifty-two, since the genotype is doubled in the phenotype. Nevertheless, this was by far the largest network we tried throughout our research. This original network also had six hidden nodes and two output nodes, governing the velocity of the hip and knee of the corresponding leg.

In the experiments on standing, the creature's legs were intentionally folded alongside its body before each trial began. Likewise, with the experiments in walking, each trial began with the legs extended straight down, already standing up. We chose to forgo the complexity of attempting to evolve the radical shift in behavior that would be necessary in order for the creature to first stand up and then transition to a walking. How to accomplish such a transition is a very interesting question definitely worthy of future research, but for this research we were not sure if we could even get walking to evolve at all, and so chose to bypass this problem of behavior transition entirely. Even with this simplification, there still exists a transitional period which must be properly handled by the set of neural networks. This results from the fact that the initial standing position is virtually guaranteed not to be a moment in the evolved walking behavior's leg-cycle. Therefore, the network must first organize the standing legs into a walking position before it can proceed with a cyclic walking behavior.

One issue that emerged was that of how to jumpstart the evolutionary process. Assuming that the initial generation consists of randomly generated genomes, which correspond to weights in the neural networks and the parameters of the oscillator signal, it is virtually guaranteed that all individuals in the initial population will fall to the ground and make no forward progress, resulting in a score of zero. The solution to this problem was found in (Mojon). They were using sin functions to govern a walking behavior and their solution to this problem was to constrain the initial population of sin functions to very low amplitudes. We adapted this idea to our neural network approach by constraining the weights impinging on the output layer to low absolute values. Since the nodes' activation functions are all sigmoidal, the output will have low absolute value when the input has low absolute value, and since the output directly effects the velocities of the joints, the joints will move in a weak fashion in the initial population, allowing for incremental improvement over evolutionary time.

As it turned out, the evolution of a basic amble gait was extremely easy. Our very first attempt to evolve walking produced a well-balanced methodic walker that, although somewhat lethargic in its motions, actually moves across the ground with tremendous forcefulness (figure 3). By this we mean that the observed walking gait is slow, but appears to be hefting the creature along as fast as it possibly can, swinging the creature's full weight heavily into every step, probably an indication that the initial joint strengths we assigned were too weak. This possibility is discussed later in the paper. A few repeated runs proved that our system could evolve good solid walkers on most evolutionary runs.

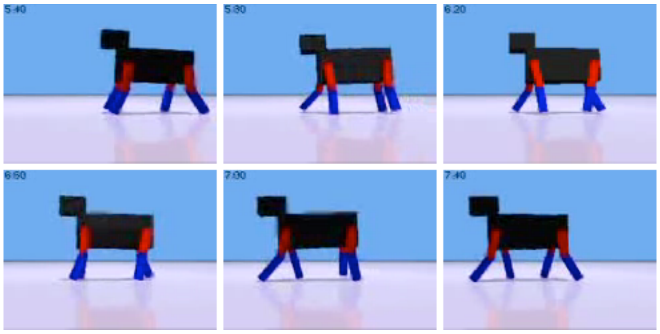

Figure 3: An example of an evolved walker demonstrating an amble gait

The study of the evolutionary progress of particular runs yielded some interesting observations. One rather surprising observation was that while there is remarkable progress made on the external performance of the walker's physical body, there are generally not very many cases of extreme change occurring within the genomes over time. After the first few generations have washed out the random clutter from the initial generation, a few founders are established, and from that point forward, the genes do not tend to vary a whole lot. Most the change seems to correspond to minor tweaking of the genes. Apparently, only small detailed changes are necessary to account for the explosive improvement in walker performance spanning most of the duration of an evolutionary run after the first few generations have passed. This it not only true of the neural network weights, which comprise the vast majority of the genome, but is also true of the three oscillator genes for the period and the left/right and front/rear phase offsets. While we were able to imagine that there very well might be some property of the neural networks that prevented them from changing in radical ways without resulting in the instantaneous death of the phenotype, we were quite surprised that the oscillator genes were so resistant to change.

It was commented earlier that it is inappropriate to think of the positive and negative phases of the oscillator signal as correlating to the physical forward swing and back swing of the actual leg. In fact, this is only partly true. Careful observation of the position of the leg at a given time reveals that it generally lags behind the oscillator signal. That is to say that the points in the leg cycle where a leg switches direction generally occur after the oscillator signal has already switched signs. Our initial theory was that either the +1s or the -1s would hold the leg in a forward or backward motion and that the sudden change in the sign of the oscillator signal would be directly responsible for the reversal of direction of a leg. We still believe that this is true, but there is a delay between the two transitions. Furthermore, the delay is more pronounced in the physical legs than in the output signals of the neural networks themselves (Figure 4). The outputs sent as commands to the legs lag slightly behind the change in the oscillator input, while the physical positions of the legs lag behind the commands they are receiving from the outputs. We believe the explanation fairly straightforward. The neural networks probably take a short period of time to shift to a new internal state when the inputs change to a new state. Furthermore, the legs are a physical structure with mass and momentum. Even after the physical legs receive commands to change velocity, the ODE physics engine does such a nice job of accurately modeling the physical system that it takes the legs a little while to fully react to the signal they are receiving and reverse direction.

We therefore conclude that our initial theory is correct. The sign transitions of the oscillator signal do indeed map directly to the forward/backward transitions of the leg itself, with the caveat that there are at least two sources of delay in the reaction of the leg to the oscillator signal. Finally, it is important to note that the sign of the oscillator signal apparently has absolutely no correlation to the direction of the leg across many trials. Each evolutionary run apparently uses the sign of the oscillator signal arbitrary as signifying the beginning of a forward or a backward leg motion. In fact, since the internal structure of the two neural networks in the genotype are entirely independent, it is quite common for the sign/direction correlation between the oscillator signal and the legs to differ between the front and rear legs.

Another interesting observation is that when we took a fully evolved walker and explicitly modified its oscillator period, the creature became instantly incapable of any successful walking at all. We were unsure whether modifying the oscillator period might simply alter the speed of the walker. In hindsight, it is fairly obvious that this should not be case. If speeding up the walker were a simple matter of modifying the oscillator period, evolution would surely have discovered this advantage through mutation of the oscillator period gene.

Therefore, it can be assumed based on the results of our explicit modification of the oscillator period, which resulted in the effective death of the walker, that the oscillator genes and the neural networks are tightly coupled, so much so, that we now believe we understand why the oscillator genes tend to vary extremely little over the course of an evolutionary run. Any evolutionary modification of these genes, even early in a run, is probably detrimental to the performance of the walker. Therefore, the success of any given single run is probably dependent on chance in the early generations, as to which individuals become dominant, and correspondingly which individuals' oscillator genes will become more or less concrete for the remainder of the run.

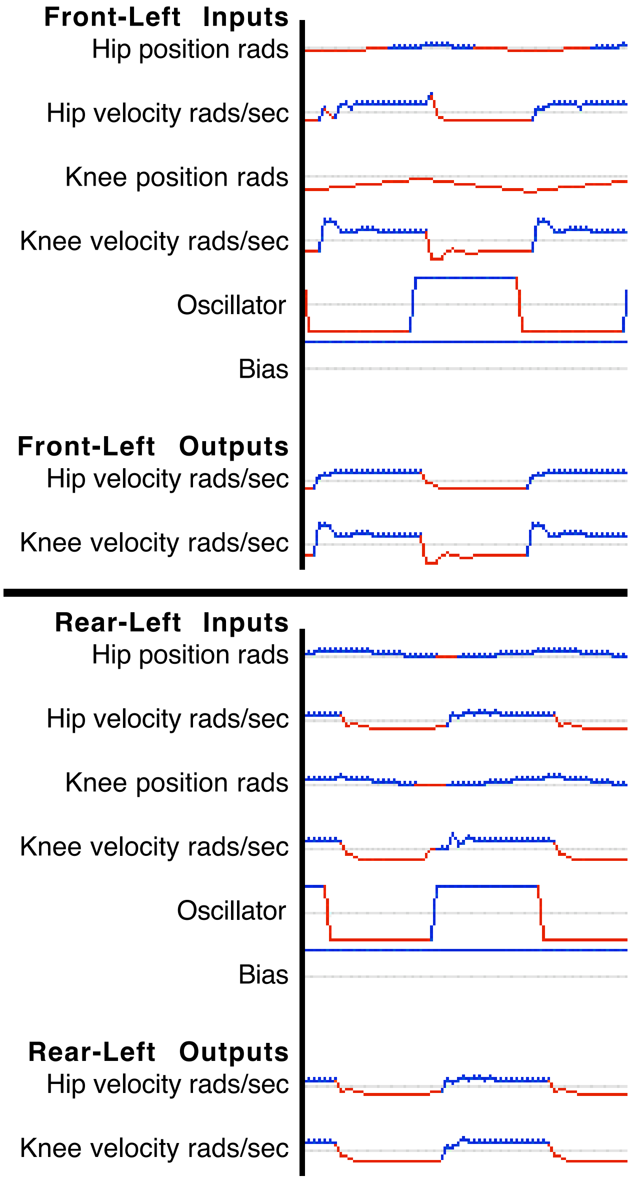

Figure 4: Neural network traces for the front-left and rear-left legs of an evolved walker (four hidden nodes in each network have been omitted to save space). Since the inputs are an instantaneous reflection of the physical legs, they can be interpreted as the actual physical qualities of the legs. Note that the phases of the neural outputs lag slightly behind the oscillator input and that the phases of the inputs lag even further behind. Admittedly, the effect is subtle. Also note that the correspondence between the sign of the oscillator and the sign of the outputs is flipped between the front and rear legs.

Minimizing the Set of Neural Inputs

Minimizing the required set of neural inputs quickly became a priority to us for two reasons, both pertaining to speed. The first is that to some miniscule degree, large neural networks take longer to calculate. This barely factored into our simulation however, where the ODE physics calculations dwarfed even the visual rendering challenges. The other problem with a large neural network however, is that it takes a much larger number of generations to search the larger genetic space for the optimal set of neural parameters. As we experimented with reducing the set of inputs (and correspondingly the number of hidden nodes as well), we discovered that we could vastly reduce the required evolutionary time in generations without sacrificing any final performance. Essentially, we were able to reduce the number of inputs for each network from the initial thirteen to only six (the major difference being that each leg no longer received the joint positions of the other legs), and at the same time reduced the number of hidden nodes from six to four, reducing the total number of neural weights in the genome from 180 to sixty-four.

Varying Floor Friction

Observations of the evolutionary progress of a walker demonstrated that the evolution progressed through a number of distinct phases. Bearing in mind that the initial population is constrained to joint velocities of low magnitude, the following description fits most runs:

1. The body shifts forward and backward but the feet do not move
2. As the body becomes more energetic in its forward/backward swings the front feet start to jitter forward slightly
3. It gets to a point where the front feet are moving the body forward a reasonable distance by shifting and sliding across the ground (without lifting into the air), all the while dragging the rear feet behind
4. The rear feet begin to get in on the action, similarly shifting along the ground
5. The feet start lifting from the ground and the body sways from side to side throughout the walking cycle
6. Solid walking is achieved in which the feet are leaving the ground and the body's weight is swinging helpfully to heft the entire creature forward at the maximum possible speed

These observations brought to light the question as to whether our system would be of much use on rougher ground, since the early stages are so utterly dependent on the ability to scrape across the ground. Rather than experiment with rough terrain, which would be a good avenue to pursue for future work, we chose to experiment with altering the friction of the flat floor. A suite of runs was performed over a range of floor frictions to see how the performance would vary with this parameter. At least two questions are of interest concerning this issue. The first corresponds to the set of observations listed above: how does the evolutionary progress over time vary with floor friction? The second concerns the final attainable speed on given floor frictions, which can be used to normalize the progress of our various experiments (figure 5). Apparently, floor friction has very little outcome on the final speed, except at very high friction values. We were quite surprised by this, as we expected that there would be a sweet spot and that lower and higher values would both demonstrate a decrease in performance. There is some differentiation in the walking styles across frictions however. We did not have time to thoroughly measure the differences but it appears that the evolved walkers at lower frictions take advantage of the ability to slide or drag their feet across the floor while the walkers at higher frictions have to pick their feet up and place them back down without dragging or sliding them. This result is fairly obvious, but it is still interesting to see the different walking styles emerge.

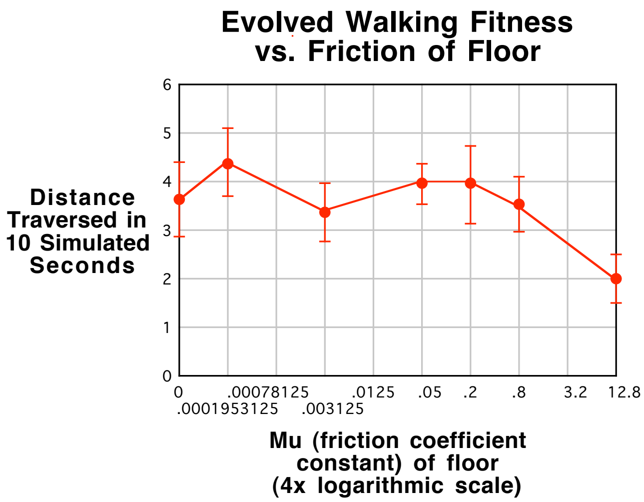

Figure 5: Final evolved fitness across a range of frictions. The default friction for most of the experiments was .2. Each point represents eight trials with a 95% confidence interval.

Varying Floor Pitch

Breve's utilities for modeling rough ground were still in a state of development at the time of this research and only become feasible toward the end of our experimentation. However we wanted to test our system on some form of terrain that presented a more difficult challenge than a perfectly flat surface, so we ran tests in which the pitch of the floor was varied (we also looked briefly into rolling the floor sideways with relation to the creature's starting position, but we did not have time to obtain conclusive results under such conditions).

A suite of runs was performed in which the pitch was varied over a range of YX slopes between -.3 (down hill) to .2 (up hill). These values represent the approximate extreme values of pitch past which the initial population would all fall over down the hill, although we did not exhaustively test the sharpness of this cutoff. It is possible that it fades out to some degree.

It should be kept in mind that the two extremes mentioned do not necessarily represent the extremes past which is impossible for the morphology we chose to traverse the hill. It is possible that an existing walker could successfully traverse steeper slopes but that evolution could not keep the initial population alive under such detrimental circumstances. We were unsure how to test this possibility as walkers that are evolved on a particular slope do not necessarily perform well on alternate slopes.

We were fairly surprised to discover that there is only a weak trend toward slower final attainable speed up hill than down hill (figure 6). Of greater interest is that the gaits are different at the more extreme pitches. On steep down hill slopes the creature extends its front legs, knees straight, in front of its body and only walks with its hind legs, push its front legs in front if itself like a wheel barrow. Likewise, on steep up hill slopes, the creature occasionally does the exact opposite, stretching its hind legs out behind itself at an angle and hauling its weight up the hill solely with its front legs, dragging its rear legs behind, although uphill gaits also emerge which more accurately represent normal walking gaits. Additionally, the only observed case of a true pace gait (in which the diagonally opposing legs move in perfect synchrony) emerged on a slight uphill incline. This was not observed on flat ground throughout our research although the similarity between the amble and pace gaits is undeniable. Many of the emergent amble gaits were almost pace gaits.

It remains to be seen whether the leg-pushing and leg-dragging strategies would remain dominant under rough terrain circumstances, where it seems intuitive that such strategies would not perform very well. Unfortunately, Breve did not allow for the variation of floor friction on the sloped floor (the sloped floor was defined as a different kind of object from the flat floor in Breve, one which did not sport a parameter for friction control), so we were unable to repeat the friction tests against the pitch tests.

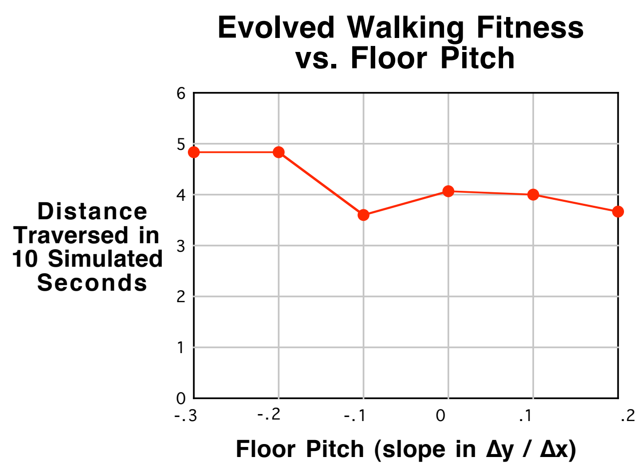

Figure 6: Final evolved fitness across a range of pitches. Although there appears to be a weak inverse correlation between pitch and fitness, this cannot be known for sure without conducting more thorough experimentation. There

are presently no error bars in the graph because each point unfortunately represents only a single run.

Varying Muscle Strength

We found that under the original conditions described earlier for the evolution of basic walking on unmodified terrain that we did not observe a large variation in the emerging gaits. While there may have been multiple causes of this result, we suspected that our original model simply had joints that were too weak, joints there were incapable of moving the legs quickly enough for more energetic gaits. Under such circumstances, it is obvious that slow gaits should emerge. In an attempt to get more energetic gaits we performed a suite of runs in which we varied the joint strength.

The results were somewhat encouraging, although not revolutionary. With higher strengths we observed the increased frequency of a nonambling gaits. In particular, a bounding gait (figure 7), or mixed variants of bounding and galloping, appear to be slightly more common at higher strengths than at lower strengths (although bounding was observed rarely at lower strengths as well). The resulting bounding gaits can be extremely energetic, and show the creature to be throwing itself around quite wildly as it takes full advantage of its greater strength. While the bounding gait does emerge at lower strengths on occasion, the creature lacks the necessary strength to actually leave the ground with all four legs at any point during the gait. At higher strengths, this does occur occasionally. Ironically, the bounding gaits do not generally outperform the methodical amble gaits however, which make steady forward progress despite their apparent attention to detail. In contrast, the bounding gaits appear to be difficult to control, such that the creature has difficulty maintaining a straight line trajectory. Since the fitness is purely a function of the X component of the final position of the creature, bounders generally score considerably lower final fitnesses with relation to the amblers.

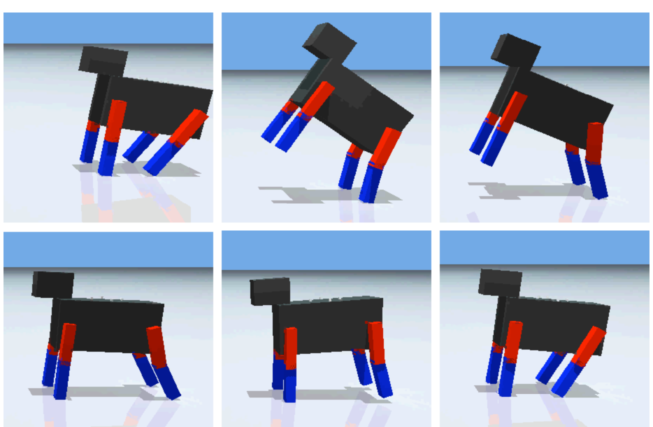

Figure 7: An example of an evolved bounding gait

One might wonder why the bounding gait ever emerges in such a case. Time analysis of the genome — the oscillator signal phase offsets in particular — suggest that a strong founder effect is taking place. All three of the oscillator signal genes tend to not vary all that much over evolutionary time. However, this raises a follow-up question. Why don't bounding gaits emerge at lower strengths more often? While it may be true that the creature lacks the necessary strength to throw itself into the air, it ought to be equally possible for the weaker bound-like gait described above. We are unsure why this result occurs more rarely at lower strengths. One possibility is that it that the body may be too well balanced or that a perfectly flat floor is too even such that it is difficult to slide or kick the front legs forward without simultaneously sliding and kicking the rear legs backward, and vs/va, in such a fashion that no forward progress can be made. A few cursory observations do seem to support this assertion, although there is insufficient evidence to make a strong claim that this is the primary explanation for the lack of emergence of bounding gaits at lower strengths. In the cases where bounding emerges at lower strengths, the creature has figured out how to lift itself up on its hind legs and then effectively fall on its front legs outstretched, and then slide the rear legs up under its body. One reason it may be possible to slide the rear legs up, but not slide the front legs out, is that the heavy head sits over the front legs, pushing them to the ground with greater force than the rear legs. Careful observation suggests that the creature is exploiting this fact since the front legs do slide backward while the rear legs are sliding forward under the body, but don't slide as far as the rear legs, and then the creature rears up on its rear legs and kicks out to fall on its front legs again.

Conclusions

We described an evolutionary system consisting of a population of neural networks and associated oscillators that are evolved to successfully control a simulated quadruped creature so as to locomote the creature across flat ground with impressive speed and elegance of gait. The breakthrough came early in the research when we discovered that doubling the genotype in the phenotype allowed us to go from being unable to evolve even simple standing behavior to being able to evolve complex walking behavior on practically every run.

Following that success, we proceeded to test our system under a variety of circumstances, including the simplification of the neural network architecture, the variation of ground friction, the variation of ground pitch, and the variation of joint strength. Perhaps the most telling result is that our system is impressively robust to a wide range of these variables, being capable of producing walkers under most circumstances, and illustrating interesting boundary cases, such as strange, yet perfectly functional, gaits when traversing steep hills.

Further, we observed the differentiation in gait styles across a range of floor frictions. On lower friction floors the gaits that emerge take advantage of the ability to slide the feet across the floor, while on higher friction floors the gaits must meticulously lift and lower their feet without sliding them across the floor. Additionally, we observed the emergence of a wider diversity of energetic gaits by permitting the creature to use stronger joints.

Future Work

Besides the obvious necessary extension to physical reality, which must be an absolute requirement if there is to be any hope of extending this research to commercial robotics, there are number of issues left which garner further experimentation. The system really needs to be tested on rough ground. Even indoor environments have thresholds between different kinds of floor and carpet.

This system has no specific notion of conscious control over the motion, including any way to specify a desire to turn. Similarly, the ability of the neural networks to differentiate between different phases of behavior, such as standing and walking, or walking and turning, have not been addressed.

To bring the above suggestions together, there should also be an emphasis on transitioning between different kinds of terrain, such as floor and carpet, instead of merely evolving for a single type of terrain and effectively dieing the instant the terrain changes.

We performed virtually no experimentation on varying the rather arbitrary morphology we originally designed. Since this work is oriented toward developing walking behaviors for future robots, and not only for existing robots, it would serve the end goal quite well to attempt to design, perhaps through evolutionary means, types of morphologies that are more suited to adept walking in the first place. The addition of ankle joints and increased degrees of freedom in the hip are obvious necessities, and we theorize that a twisting or turning torso would probably increase the creature's agility as well, based purely on our own observations of existing successful quadrupeds such as cats and dogs.

Lastly, we would suggest that the inclusion of true neural CPGs would make a fine addition to this work, since it would allow for the entire behavior to be encompassed in a neural network design without the need or externally imposed oscillator signals.

Acknowledgements

Throughout this research, a number of conversations conducted with Dr. David Ackley (University of New Mexico, Albuquerque, NM) proved quite helpful in guiding the thought process and experimental design.

References

Barnes D. 1998. Hexapodal Robot Locomotion Over Uneven Terrain. *Proc. IEEE Conf. on Control Applications*. pp. 441 - 445.

Chiel J, Beer, R, Gallacher, J. 1999. Evolution and Analysis of Model CPGs for Walking I. Dynamical Modules. *Journal of Computational Neuroscience* 7:99-118.

Golubitsky M, Stewart I, Buono P, and Collins J. 1999. Symmetry in Locomotor Central Pattern Generators and Animal Gaits. *Nature*. 1999. 693-5.

Hornby G, Fujita M, Takamura S, Yamamoto T, and Hanagata O. 1999. Autonomous Evolution of Gaits with the Sony Quadruped Robot. *Proceedings of 1999 Genetic and Evolutionary Computation Conference (GECCO).* Banzhaf, Daida, Eiben, Garzon, Honavar, Jakiela, Smith, eds., Morgan Kauffmann, pp. 1297-1304.

Hornby G, Takamura S, Hanagata O, Fujita M, and Pollack J. 2000. Evolution of Controllers from a High-Level Simulator to a High DOF Robot. Miller, J. (ed) *Evolvable Systems: from biology to hardware; proceedings of the third international conference (ICES 2000).* Springer (Lecture Notes in Computer Science; Vol. 1801). pp. 80-89.

Hornby G, Takamura S, Yokono J, Hanagata O, Yamamoto T, and Fujita M. 2000. Evolving Robust Gaits with AIBO. *IEEE International Conference on Consumer Electronics*. pp. 192-193.

Ijspeert A. and Kodjabachian J. 1999. Evolution and Development of a Central Pattern Generator for the Swimming of a Lamprey. *Artificial Life* 5:3, pp 247-269.

Klein J. 2002. BREVE: a 3D Simulation Environment for the Simulation of Decentralized Systems and Artificial Life. *Proceedings of Artificial Life VIII, the 8th International C*onference on the Simulation and Synthesis of Living Systems. p. 329.

Lewis M. 2002. Gait Adaptation in a Quadruped Robot. *Autonomous Robots* 12. (3): 301-312.

Mojon S. Realization of a Physic Simulation for a Biped Robot. Not published.