Constructing Controllers for Physical Multilegged Robots using the ENSO Neuroevolution Approach

Vinod K. Valsalam[1], Jonathan Hiller[2], Robert MacCurdy[2], Hod Lipson[2], and Risto Miikkulainen[1]

[1]Department of Computer Sciences, The University of Texas at Austin, Austin, TX 78712 USA
[2]Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

Abstract

Evolving controllers for multilegged robots in simulation is convenient and flexible, making it possible to prototype ideas rapidly. However, transferring the resulting controllers to physical robots is challenging because it is difficult to simulate real-world complexities with sufficient accuracy. This paper bridges this gap by utilizing the Evolution of Network Symmetry and mOdularity (ENSO) approach to evolve modular neural network controllers that are robust to discrepancies between simulation and reality. This approach was evaluated by building a physical quadruped robot and by evolving controllers for it in simulation. An approximate model of the robot and its environment was built in a physical simulation and uncertainties in the real world were modeled as noise. The resulting controllers produced well-synchronized trot gaits when they were transferred to the physical robot, even on different walking surfaces. In contrast to a hand-designed PID controller, the evolved controllers also generalized well to changes in experimental conditions such as loss of voltage and were more robust against faults such as loss of a leg, making them strong candidates for real-world applications.

1 Introduction

Imagine having to design a multilegged robot to explore Mars. Such a robot could navigate the rugged terrains of Mars better than a wheeled robot, but it is more difficult to control. The controller must coordinate the robot's legs properly, generating robust gaits to negotiate different terrains effectively while maintaining stability. Moreover, the robot should be robust to different environmental conditions, wear and tear, and even failure like losing one or more legs, to reliably complete its mission.

Research on designing controllers for multilegged robots began in the 1980s [1, 6, 29] and continues actively both in academia and industry [3, 8, 15, 19]. The controllers for state-of-the-art robots are often designed by hand, requiring extensive analysis of the sensor-motor systems and body-limb dynamics [3, 8, 15, 19, 29]. This process is generally difficult and brittle because it is hard to anticipate all operating conditions.

Therefore, automatic design methods utilizing learning techniques such as evolution are a promising alternative [2, 4, 14, 35]. These methods typically evaluate controller fitness by simulating the physical behavior of robots and their environments. Such experiments are useful because they allow testing a range of conditions quickly and effectively without damaging expensive hardware components of the robot. However, simulation may not always produce accurate enough results, making it difficult to transfer the evolved controller to the physical robot [22, 23].

This paper evaluates the hypothesis that controllers evolved in an approximate simulation can be transferred successfully to the physical robot if the evolved controllers are robust enough. An approximate simulation is first made by modeling the weight distribution of the robot and the uncertainties in its sensors and actuators. An approach called Evolution of Network Symmetry and mOdularity [ENSO; 33, 35, 36] is then utilized to evolve controllers that transfer robustly from simulation to reality. This method was demonstrated on a physical quadruped robot that was designed and built using rapid prototyping technology. The evolved controllers produced the same gaits both in simulation and on the physical robot. Moreover, they generalized better to different experimental conditions and demonstrated better fault tolerance than a hand-designed PID controller, suggesting that such evolved controllers are useful for real-world applications.

2 Background

This section reviews prior work on evolving controllers for physical robots and summarizes the ENSO approach that this paper utilizes to evolve controllers.

2.1 Evolving Controllers for Physical Robots

It is possible to evolve controllers by evaluating their fitness directly on the real robot instead of in simulation [11, 16, 20, 37, 40]. However, performing thousands of such fitness evaluations in hardware may be impractical for several reasons [23]. First, hardware evaluations are slow, resulting in long evolutionary run times. Second, they cause wear and tear on the robot, making hardware failures likely and user intervention to repair them necessary. Third, the controllers created by evolution through random variations may produce abnormal actuator signals that can crash or damage the robot.

Therefore, a good alternative is to evaluate controller fitness in simulation and then transfer only the final, evolved controller to the physical robot. However, transferring such controllers evolved in simulation to the physical robot is challenging [7, 22, 23]. The main reason is that it is difficult to simulate physical properties such as friction and sensor and

actuator characteristics with high enough fidelity to reproduce the simulated behaviors on real robots. In order to address this issue, researchers have developed several methods that improve the results of evolution in simulation by performing only a few experiments on the real robot.

A straightforward method is to fine-tune the controller behaviors by continuing evolution on the real robot for a few more generations [24, 25]. However, this method may be ineffective in correcting behaviors that have evolved to exploit flaws in the simulation. A better alternative is to make such behaviors less likely to evolve by incorporating transfer experiments from the beginning of evolution, e.g. by utilizing a multi-objective evolutionary algorithm that optimizes both a task-dependent controller fitness as well as a measure of how well the controller transfers from simulation to reality [21]. In any given generation, this method chooses at most one controller based on behavioral diversity to be evaluated on the real robot, requiring only a small number of hardware evaluations.

Other methods utilize the information they gather from a few experiments on the real robot to build a more realistic simulator, typically in one of two ways: (1) Experiments are performed on the real robot before running evolution to collect samples of the real world by recording sensor activations [24, 25]. When controllers are evaluated later during evolution, these samples are utilized to set the simulated sensor activations accurately. (2) Experiments are performed on the real robot during evolution to co-evolve the simulator and the controller, making an initially crude simulation more and more accurate [5, 7, 38].

In contrast to the above methods, this paper proposes to bridge the simulation-reality gap simply by creating controllers that are robust to small discrepancies between simulation and reality. It utilizes the ENSO approach [33, 35, 36], which evolves controllers modeled as coupled cell systems to provide theoretical guarantees of robustness. Therefore, simulation accuracies sufficient for ENSO can be obtained easily by approximating the morphology and mass distribution of the real robot with cylindrical and rectangular blocks, without performing any hardware experiments either before or during evolution (Section 3.2).

Evolving controllers in an accurate enough simulation is often insufficient to transfer them successfully to the real world because of uncertainties in sensor activations and actuator responses. However, past work has shown that evolution can adapt controllers to such uncertainties by modeling them as noise in the simulation [13, 17, 18, 24]. The same idea is utilized in this paper, focusing on uncertainties in how the motors respond to control signals. The resulting controllers transfer successfully to the physical robot, producing the same behaviors both in simulation and on the physical robot. The ENSO approach that is used to evolve these controllers is discussed next.


Figure 1: **Example symmetries of a modular neural network controller for a quadruped robot.** The four controller modules and their interconnections can be represented as the vertices and edges of a colored graph. Vertices and edges of the same color (indicated by line style) are identical, each color representing a particular combination of neural network parameters. A graph symmetry is any permutation of vertices under which the edge colors remain the same. Both graphs in this figure have vertices of the same color. Moreover, all edges of graph G_A have the same color, while edges of graph G_B have different colors. Therefore, any permutation of the vertices of graph G_A is a symmetry. In contrast, only the permutations $(1\ 2)(3\ 4)$ and $(1\ 3)(2\ 4)$, and their compositions are symmetries of graph G_B. The set of all symmetries of a graph forms a group. For example, the symmetry group of graph G_A is the *symmetric group* $\mathcal{S}_4$ and that of the less symmetric graph G_B is a subgroup of $\mathcal{S}_4$ called the *dihedral group* $\mathcal{D}_2$. The symmetry group of the controller determines the type of gaits that it can produce [9].

2.2 The ENSO Approach

The controller for a multilegged robot can be implemented as a system of interconnected neural network modules, each controlling a different leg [1, 34]. Some of these modules and interconnections may be identical, resulting in symmetries, i.e. permutations of the modules that leave their interconnection graph invariant (Figure 1). Symmetries can be analyzed using group theory to show that they determine the type of gaits that the controller can produce [9]. However, designing the appropriate symmetries by hand is often difficult and may even be infeasible in the general case [35]. Moreover, the parameters of the neural network have to be optimized together with its symmetries. The ENSO approach was designed to evolve solutions to such problems automatically [33, 35, 36], and it is therefore utilized in this paper.

ENSO initializes evolution with a population of maximally symmetric solutions (Figure 2): They have the simplest possible structure, consisting of identical modules and interconnections. This structure is represented by a colored graph with identical vertices and edges. During evolution, group-theoretic mutations [33, 35, 36] break their symmetries systematically, thus exploring less symmetric, more complex solutions with different types of modules and interconnections. Each such mutation creates a new graph with more colors such that its symmetry group is a random maximal subgroup of the original symmetry group.

ENSO organizes the colors created by successive symmetry (color) mutations as a tree. Each such tree is a genotype for evolution. The leaf nodes of this tree represent the colors of identical vertices and edges of the phenotype graph, which in turn represents a neural network (Fig. 2a). In particular, each leaf node specifies which vertices or edges have that color and stores the corresponding set of neural network parameters, i.e.
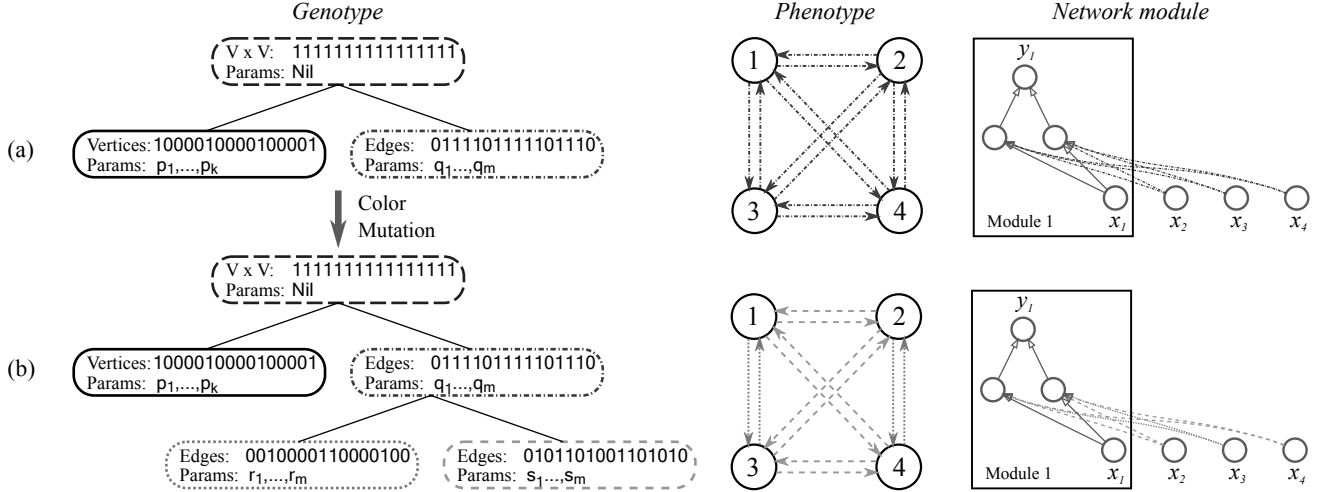
Figure 2: Examples of genotype, phenotype, network module, and color mutation. ENSO uses a tree of colors as genotype (left). Each leaf of this tree has a unique color, and represents a set of vertices or edges of the phenotype graph (middle) that have the same parameter values. The vertices and edges of the phenotype graph represent the modules of a neural network and the connections between them (right). Their parameters (stored in the genotype) consist of node biases and connection weights for each module network (vertex) and weights for each connection between modules (edge). Each module has a fixed architecture with a layer of hidden nodes fully connected to its inputs and outputs. A connection from another module (not shown) is implemented by fully connecting its input layer to the hidden layer of the target module. (a) At the beginning of evolution, each genotype in the population represents a maximally symmetric phenotype graph with symmetry group $\mathcal{S}_4$. All vertices of this graph have the same color (solid, represented by the leaf on the left) and all its edges have the same color (alternating dots and dashes, represented by the leaf on the right), implying that all modules are identical and all connections between them are also identical. (b) A color mutation breaks the phenotype graph symmetry to $\mathcal{D}_4$, which is a maximal subgroup of $\mathcal{S}_4$. As a result, two child nodes are created for the node representing the set of edges, i.e. the set of edges is partitioned into two and each part is colored differently (dotted and dashed). Since each color is associated with a different combination of parameter values, the mutated phenotype graph represents two types of connections between network modules. Such color mutations, when combined with parameter mutations, make it possible to evolve symmetric and modular neural networks efficiently.

the biases and connection weights of the module network (for vertex color) or the connection weights between modules (for edge color).

Symmetry mutations produce structural changes in the phenotype graph by partitioning the vertex or edge set of leaf nodes and creating a new child color for each part of the partition (Fig. 2b). Structural changes are also produced by enabling or disabling edges of the same color using edge-toggle mutations. Interleaving these two types of structural mutations with parameter mutations optimizes the modules and interconnections of simpler solutions first and elaborates on them to create more complex solutions only when necessary. Moreover, the systematic, symmetry-breaking mutations constrain the search to promising symmetries, making evolution more effective than mutating symmetry unsystematically [33, 35, 36].

These features make it possible for ENSO to evolve complex modular systems such as the quadruped controller illustrated in Figure 3. All four modules of this controller utilize the same two-layered neural network architecture (Figure 3a). Each module's input is the joint angle of the leg it controls. It can be represented by the angle itself, or by the sine and cosine of the angle; the sine and cosine are actually more robust (because they are continuous), and were used to obtain the results discussed in Section 4. The module's output is the desired angular velocity of that leg. The full controller network is obtained by connecting the four modules to each other

such that each module receives input from all the other modules (Figure 3b).

The phenotype graph representing the symmetries of the controller network also represents a central pattern generator, or coupled cell system, i.e. a dynamical system capable of generating robust periodic oscillations for modeling quadruped gaits [9]. Therefore, ENSO can design effective controllers by evolving their symmetries in addition to their parameters, as demonstrated previously in physically realistic simulations of a quadruped [35, 36]. Symmetry evolution was especially advantageous when the appropriate symmetries are difficult to determine manually, such as on an inclined ground. In particular, ENSO evolved gaits that were significantly faster and also generalized better than those evolved with only parameter mutations on hand-designed symmetries. Moreover, the group-theoretic symmetry mutations of ENSO produced smoother and better coordinated gaits than those produced using arbitrary symmetry mutations. Since the ability to evolve such robust and regular gaits is important for real-world applications, ENSO is a promising approach for designing distributed controllers for real robots. This paper investigates this hypothesis by extending previous simulation results, as will be described next.

3

(a) Network module and the initial phenotype graph

(b) Full controller network consisting of four modules

Figure 3: Modular controller network for a quadruped robot. The input to each module is the angle (or its sine and cosine) of the leg it controls, and the output is the desired angular velocity of that leg. The full controller network consists of four such modules, each module receiving input from all the other modules. The symmetries of these modules and their connectivity are represented by a phenotype graph. At the beginning of evolution, this graph has identical vertices (modules) and edges (interconnections), i.e. all vertices and edges have the same combination of network parameters. ENSO discovers effective controllers by breaking symmetry to create new types of vertices and edges, and by optimizing their initially random network parameters.

3 Method

A physical quadruped robot was designed and fabricated to evaluate the ENSO approach in real-world applications. This section discusses its design, simulation, and control.

3.1 Physical Robot

The physical robot was designed with two constraints: (1) it must be similar to the simulated robots that were previously studied using ENSO [35, 36], and (2) it must be possible to prototype it quickly with parts that can be purchased or fabricated easily. While some parts of the robot such as the servo motors and the controller board to operate the motors were available commercially, other parts such as its body and legs were custom-designed and fabricated to fit the commercial parts.

Each leg of the robot is attached to a Dynamixel AX-12+ servo motor manufactured by Robotis [30]. The AX-12+ provides angular position feedback, and can be made to rotate continuously by specifying the desired angular velocity, thus matching the inputs and outputs of the neural network controllers that ENSO evolves. The evolved controller runs on a Robotis CM-2+ microcontroller circuit board, which provides an interface to communicate with the Dynamixel motors through a daisy-chain serial connection.

The AX-12+ motors are mounted on a rectangular body using a wedge-shaped piece to tilt their axes of rotation 20° from the vertical (Figure 4). The legs also slant 20° from their re-


Figure 4: Assembled physical quadruped robot. The Dynamixel AX-12+ motors rotate the legs of the robot and are mounted on the four corners of its rectangular acrylic body, which has attachment points in the middle for a future hexapod extension. The legs make 20° with the axis of rotation of their respective motors, tracing cones as they rotate. The motor axes also have a 20° sideways tilt from the vertical. As a result, rotating the legs raises and lowers them and can produce locomotion when they make contact with the ground. A control program synchronizes the rotation of the legs to produce locomotion and runs on the CM-2+ circuit board mounted on top of the body. The board is powered by a 12V lithium-ion battery attached to the under-side of the body. Videos of experiments using this robot can be seen at the website http://nn.cs.utexas.edu/?enso-realrobots

spective motor axes, making it possible for the robot to walk by rotating its legs continuously. The body, the wedge, and the leg were designed using SolidWorks [31], a program for computer aided design. The body was then cut from acrylic using a laser cutter and the wedges and the legs were fabricated in an Objet Eden 260V [26] rapid-prototyping 3D printer. The circuit board is mounted on the top side of the body and is powered by a 12V lithium-ion battery attached to its bottom side by Velcro.

This robot was then modeled in a physics simulation for evolving neural network controllers using ENSO. The simulation utilized OPAL [28], an abstraction library on top of the Open Dynamics Engine (ODE) [27], and is described in detail next.

3.2 Simulation

The robot's legs were modeled as cylinders with capped ends, while its body was assembled from several rectangular boxes that approximate different body parts (Figure 5). These cylinders and boxes have the same dimensions and the same relative angles as the corresponding parts in the physical robot. The leg angles used as controller input are measured from the same vertical leg positions. Moreover, densities were assigned to the parts such that they have the same mass in both simulation and the real robot. As a result, the simulation utilizes the same approximate morphology and mass distribution as the physical robot.

In addition to the morphology, the sensor and actuator characteristics of the AX-12+ motors were also modeled. The mo-

Figure 5: Simulation model of the physical quadruped robot. This model simulates the morphology and dynamics of the physical robot in Figure 4. The legs were modeled as capped cylinders and the other parts were approximated as rectangular boxes with the same dimensions. The weights of these shapes were then matched with those of the corresponding parts of the physical robot. As a result, this model represents the weight distribution of the physical robot with sufficient accuracy for simulation.
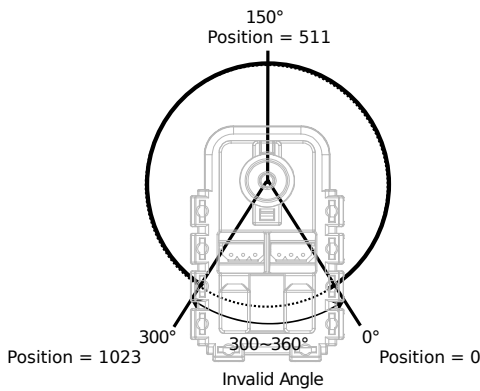

Figure 6: Angular position sensor readings of the Dynamixel AX-12+ motor. The angular position sensor of the motor provides an integer valued reading in the range $[0, 1023]$ when the motor is in the $[0, 300]$ degree range; angles outside this range produce invalid sensor readings and are therefore interpolated. Reprinted with permission from Robotis [30]; annotations edited.

tor can sense its angular position if it is in the $[0, 300]$ degree range (Figure 6). However, it does not give valid position feedback for angles between $300°$ and $360°$. Since the neural network controllers take angular positions as inputs, the sensor reading is interpolated when the motor is in this blind zone. In fact, the sensor reading is calculated the same way for all angular positions from an estimate of the angular velocity, which gets updated only when the angle is in the valid range. Exponential smoothing is applied to this estimate to filter out noise and discontinuities caused by any discrepancy between the estimated and actual angular velocities when the motor emerges from the blind zone.

The response of the motor to the angular velocity control signals from the neural network is more difficult to model accurately. In particular, the angular velocity of the motor drifts significantly over time for a constant control signal.

This stochastic drift can change the periodic trajectory of the controller, thus disrupting the gait of the robot. Such uncertainties can be handled by adding noise to the simulation, making it possible for evolution to adapt the controllers suitably [13, 17, 22]. Two types of Gaussian noise were added: The first type models fluctuations about the mean with standard deviation 2.5%. The second type of noise models drifts in the mean; it is therefore larger in magnitude (standard deviation 20%), but it is applied only a few times in each evaluation.

The accuracy of the motor in representing its angular position and velocity is also included in the simulation. The motor represents both variables with an integer value in the range $[0, 1023]$. The precision of these variables in simulation is therefore downgraded from floating point precision to match the actual precision on this robot. Moreover, the neural network controller reads the positions and updates the velocities at the same frequency in both the simulation and the real robot.

The controller evolved in simulation in the above manner is transferred to the real robot by programming the CM-2+ circuit board with it, as described next.

3.3 Control Programs

The CM-2+ circuit board contains an ATmega128 CPU, which is an 8-bit microcontroller with 128KB of on-chip programmable flash memory. Running the programs stored in its memory can activate the motors of the robot. Therefore, the evolved neural network controller is interfaced with the motors by converting it to a C-language program and invoking it from a control loop similar to that used in simulation. This C program is then cross-compiled for the ATmega128 using the GNU compiler toolchain and downloaded to the CM-2+ board through an RS-232 serial connection.

This facility for writing control programs in C was also utilized to hand-code a baseline controller for comparing with the evolved controllers. The hand-designed controller replaces the neural network in the control loop with a PID controller that also controls the leg angular velocities utilizing feedback of leg positions. Thus, the hand-designed controller also benefits from the mechanism mentioned in Section 3.2 for interpolating and smoothing sensor readings of leg positions. It computes the error signals for PID control as the difference between the actual leg positions and the desired leg positions for obtaining uniform leg angular velocities.

The results of experiments comparing this controller with the evolved controllers are described in the next section.

4 Results

The gaits produced by the evolved and hand-designed controllers were evaluated for walking on flat ground (1) when all four legs of the robot are functional and (2) when one leg is disabled to simulate a real-world motor failure. The controllers produced in the first experiment were evaluated further for generalization by reducing the maximum speed of the motors and by initializing one of the legs with a large error. In each experiment, generalization was also tested by placing the

robot on different surfaces. These experiments and the resulting gaits are discussed in the following subsections. Videos of example gaits can be seen at `http://nn.cs.utexas.edu/?enso-realrobots`.

4.1 Experimental Setup

In each evolutionary run, the initial population of controller networks had the architecture depicted in Figure 3. Their connection weights were set randomly within $[-2, 2)$, neuron biases to 0, and neuron sigmoid slopes to 1. During evolution, these parameters were mutated with Gaussian perturbations (with $\sigma = 0.2$) acting with a specified probability (0.5). All edges were enabled in the phenotype graphs of the initial controllers, and mutations toggled them with a specified probability (0.1). In each generation, an offspring was created by first selecting a parent in a two-way tournament, and then applying either a parameter mutation, an edge-toggle mutation, or a symmetry mutation. Parameter mutations were 100 times more likely and edge-toggle mutations were ten times more likely than symmetry mutations. Each symmetry mutation created five offspring, all having the same symmetry, and their newly created parameters were initialized randomly. In addition, the network with the best fitness was copied without change to the next generation. A population size of 200 was used in all experiments. These particular settings were found to work well empirically and small variations did not produce a noticeable impact on the evolved results.

Each network was evaluated in a simulation (Section 3.2) in which the network controlled the locomotion of a robot. When the robot was initially placed in the simulation environment, its longitudinal and lateral axes were aligned with the coordinate directions of the ground plane. The simulation was then carried out for one minute of simulated time with step size 0.01s, after which the fitness of the controller network was calculated as the distance traveled by the robot along the longitudinal axis. This fitness measure ensured that the controller networks were rewarded for moving forward in a straight line.

For all experiments, evolution was run for 500 generations and repeated 10 times, each time with a different random number seed. The following subsections discuss the results of each experiment in detail.

4.2 All Legs Enabled

Figure 7 shows the phenotype graph of a champion neural network controller that ENSO evolved by utilizing the simulation model described in Section 3.2. Its symmetry group is similar to that of the graphs that ENSO evolved for a simpler simulated quadruped in prior work [35]. Therefore, it generates a similar, well-synchronized trot gait (Figure 8a), thus extending the previous results to a more complex and realistic robot model.

Moreover, transferring this controller to the physical robot reproduces the same gait accurately (Figure 8b). The robot walks smoothly in a straight line even on very different surfaces such as linoleum and carpet, demonstrating that ENSO


Figure 7: Phenotype graph of a champion neural network controller evolved by ENSO. This graph has symmetry group similar to the graphs that ENSO evolved for a simpler simulated quadruped in prior work [35]. As a result, it generates a similarly well-coordinated trot gait and it transfers well from simulation to the real robot (Figure 8).

(a) Simulated robot

(b) Real robot

Figure 8: A trot gait evolved in simulation and transferred to the real robot. The plots show the four leg angles of the robot in the first twelve seconds. They were produced by the controller with the phenotype graph illustrated in Figure 7. In both plots, after the controller reaches a steady state in about four seconds, it maintains synchrony and phase relations between the legs, producing a well-coordinated trot gait. This gait works on various surfaces robustly. Both plots are very similar, indicating that the controller produces the same walking behavior in both the simulated model and the real robot.

can evolve such robust controllers that transfer successfully from simulation to real robots.

The hand-designed PID controller was also tested on the real robot with a reference waveform for a trot gait having approximately the same period as the evolved controller (Figure 9). The legs are first positioned on the reference waveform so that there is no error for the controller to correct when it starts. Thereafter, the PID mechanism of the controller cor-
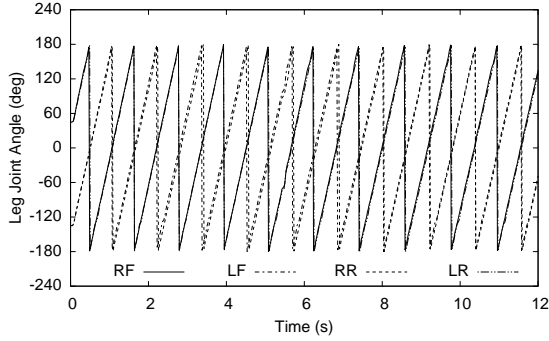
Figure 9: Trot gait produced by the hand-designed controller for the real robot. The plot shows the four leg angles of the robot in the first twelve seconds. This controller keeps the legs synchronized with a reference waveform for a trot gait by applying PID control to correct small errors in leg positions. However, it is difficult to design such controllers by hand to work robustly in the general case.

rects small errors by speeding up or slowing down the legs to keep them aligned with the reference waveform. As a result, it produces a trot gait similar to the evolved controller. However, it is not as robust and does not generalize as well as the evolved controller, as demonstrated by the experiments discussed next.

4.3 Generalization to Reduced Motor Speed

The maximum leg angular velocity that the motors can produce depends on the conditions in which the robot operates. For example, it decreases when the input voltage to the motor decreases as a result of e.g. low battery charge or temperature [12, 39]. The challenge for the controller is to keep the legs synchronized and to keep the robot walking effectively even in such conditions. The hand-designed and evolved controllers were tested for their ability to generalize to such conditions by reducing the maximum angular velocity that the motors produce.

The hand-designed controller fails, losing leg synchronization, even for a small (10%) reduction in the maximum angular velocity (Figure 10a). It fails because the legs can no longer move fast enough to keep up with the reference waveform. Slowing down the waveform can correct the problem, but doing so in a way that produces the fastest possible gait robustly is difficult, because it requires controlling both the waveform and the leg angular velocities simultaneously. In contrast, the evolved controller continues to function robustly even when the maximum angular velocity is reduced by 60% (Figure 10b). It achieves this robustness by slowing down the legs automatically and keeping them synchronized. Thus, the evolved controller generalizes well to a range of motor speeds, while the hand-designed controller generalizes poorly.

4.4 Generalization to Different Leg Positions

Another situation in which the hand-designed controller performs poorly is when the error between the position of a leg and the reference waveform becomes too large, which could happen e.g. when the leg is obstructed by an obstacle. The larger the error, the longer it takes the PID mechanism of the

**(a)** Hand-designed controller for 10% reduction in maximum motor speed

(b) Evolved controller for 60% reduction in maximum motor speed

Figure 10: Gaits produced by the hand-designed and evolved controllers on the real robot when the maximum speed of the motors is reduced. The plots show the four leg angles of the robot in the first twelve seconds. They were produced by the same hand-designed controller that produced the gait in Figure 9 and the same evolved controller that produced the gait in Figure 8. (a) Reducing the maximum speed of the motors even by 10% causes the hand-designed controller to lose leg synchronization quickly because it cannot keep up with the reference waveform. (b) The evolved controller maintains leg synchronization and performs robustly even when the maximum speed of the motors is reduced by 60%. It does so by simply slowing down the gait automatically. Thus the evolved controller is more general than the hand-designed controller.

hand-designed controller to correct the error. During this time, the leg may not be synchronized well enough with the other legs to produce a good gait. The worst such behavior occurs when the error is maximum, i.e. when the leg is 180° out-of-phase with the reference waveform.

In order to evaluate robustness against such errors, the legs were first positioned such that one leg (the left-rear one) had the maximum error of 180° and the other legs had zero error. The controller was then initialized with these leg positions, making it possible to observe how quickly it corrects the error of the left-rear leg.

The hand-designed controller takes more than 30 seconds to correct this error (Figure 11a). During this time, the angular position of the left-rear leg overshoots and undershoots the reference waveform several times, synchronizing with the right-front leg to produce the original trot gait only gradually.

(a) Hand-designed controller

(b) Evolved controller

Figure 11: Gaits produced by the hand-designed and evolved controllers on the real robot when the left-rear leg is initialized with maximum angular position error. The plots show the variation of the four leg angles of the robot with time. They were produced by the same hand-designed controller that produced the gait in Figure 9 and the same evolved controller that produced the gait in Figure 8. In order to produce the original trot gaits, the controllers must correct the initial error by synchronizing the left-rear leg with the right-front leg. (a) The hand-designed controller adjusts only the behavior of the left-rear leg to correct the error. As a result, the left-rear leg leads and trails the right-front leg alternately, eventually synchronizing only after more than 30 seconds. (b) In contrast, the evolved controller adjusts the behaviors of multiple legs simultaneously, correcting the error and achieving synchronization smoothly in about two seconds. Thus the evolved controller generalizes well, while the hand-designed controller is less robust.

Meanwhile, the other three legs that were initialized with zero error track their respective reference waveforms closely from the beginning. This behavior is the consequence of correcting the error of each leg separately without modifying the behavior of the other legs.

In contrast, the evolved controller takes only about two seconds to correct the same error (Figure 11b). Moreover, it synchronizes the legs without producing the undesirable overshooting and undershooting oscillations (ringing) that the hand-designed controller produces. This robust behavior is possible because the control module for each leg utilizes inputs from the other legs as well. As a result, the evolved controller adjusts the behavior of all legs simultaneously, bringing them into the appropriate relative phases much quicker and smoother than the hand-designed controller.


Figure 12: Phenotype graph of a champion neural network controller evolved with the left-rear leg disabled. The symmetry group of this graph is surprising in that ENSO did not evolve a different module for the disabled leg; instead, it evolved the same module for all four legs. As a result, it produces a gait resembling trot with the disabled leg not responding (Figure 13). ENSO adapted this gait to make the robot walk straight with only three legs.

Evolution utilizes this ability of one control module to influence the behavior of the other modules in the next experiment as well, producing a straight and effective gait even when one leg is disabled.

4.5 One Leg Disabled

An important requirement for many robots in the real world is fault tolerance [10]. For example, hardware failures are likely in a robot operating in hostile environments or exploring another planet. In such applications, it is not always possible to replace a failed leg actuator with a new one. Walking with the same gait as before is also not an option because the asymmetric action of the remaining legs would cause the robot to curve to one side.

In such a situation, a new controller must be designed to make the robot walk effectively with its remaining legs, recovering as much performance as possible. Designing such a controller by hand is challenging; even designing the appropriate symmetry by hand is challenging. In contrast, ENSO should be able to evolve effective neural network controllers for such an asymmetric robot automatically by disabling the failed leg in simulation. The new controller can then be downloaded to the physical robot for a successful walk.

This hypothesis was tested by evolving controllers with the left-rear leg disabled in the simulation. Figure 12 illustrates the symmetry of a resulting champion controller. Surprisingly, it is still very symmetric with the same type of module controlling all four legs. Therefore, it produces a gait similar to a trot (Figure 13) and sends activation to the disabled leg also. Since the disabled leg does not respond, evolution adapted the gait accordingly to produce a straight walk utilizing only three legs. When this controller is transferred to a similarly disabled physical robot, it produces the same gait on that robot as well, thus demonstrating successful transfer.

5 Discussion and Future Work

In the above experiments, ENSO evolved controllers that produced the same gaits in the physical robot as they did in simulation. Moreover, these gaits were robust to uncertainties that commonly occur in the real world such as changes in ground friction between different surfaces, the maximum angular velocity that the motors can produce, and the initial positions of

(a) Simulated robot

(b) Real robot

Figure 13: A gait evolved in simulation with the left-rear leg disabled and transferred to the similarly disabled real robot. The plots show the four leg angles of the robot in the first twelve seconds. They were produced by the controller with the phenotype graph illustrated in Figure 12. The disabled leg produces a flat line, while the other three legs maintain synchronous and phase-related oscillations resembling a trot gait. However, the synchronous lines for left-front and right-rear legs split slightly from each other between 0° and 180° and the line for the right-front leg curves a little around 0°, indicating adaptation of the gait to produce a straight walk with only three legs. This is a gait that transfers well to the physical disabled robot.

legs from which the robot starts walking. The transfer was successful because of two factors: (1) it was possible to simulate the physical characteristics of the robot with sufficient accuracy, and (2) ENSO was able to evolve robust controllers that produce stable gaits.

It is often difficult to transfer controllers evolved in simulation to the real world because it is difficult to simulate the physical robot and its environment accurately [13, 22]. However, this requirement is less critical for the controllers that ENSO evolves because they are coupled cell systems [9, 33, 35, 36]. Such systems are robust to small perturbations and can thus compensate for small deviations from ideal behavior. For example, instead of detailed mesh models of the robot morphology, a crude approximation of its weight distribution was sufficient. Moreover, it turned out sufficient to model the idiosyncratic properties of the motors and uncertainties in their behavior using interpolation and noise.

Modeling the controllers as coupled cell systems also makes it possible to describe the gaits they produce in terms of their symmetries [9]. Based on this observation, ENSO discovers


Figure 14: Variations in phenotype graphs of champion controllers evolved by ENSO. These controllers were evolved using different random number seeds for the experiment discussed in Section 4.2. They have different structures either with identical modules (left graph and Figure 7) or different modules for the front and rear legs (middle and right graphs). However, in all cases, the resulting symmetries were optimized by ENSO to produce effective trot gaits.

effective gaits by evolving the appropriate symmetry groups for the coupled cell systems. ENSO can discover many such roughly equivalent solutions when initialized with different random number seeds. For example, nine out of ten trials produced unique phenotype graph representations of controllers for the experiment discussed in Section 4.2. An arbitrary sample of three such controllers is illustrated in Figure 14. Although their structures are different, their symmetries were optimized by ENSO to produce similar trot gaits, resulting in roughly the same fitness. Moreover, the controller depicted in the left graph has identical modules, similar to the controller in Figure 7. In contrast, the controllers depicted in the middle and the right graphs have different modules for the front and rear legs, making it possible to tune the generated gaits if necessary by evolving their behaviors independently. Since different modules can represent different control functions, such graphs are more likely to evolve in more complex platforms that require specialized leg behaviors.

Evolving controllers in simulation and then transferring them to the real robot in this manner is an effective alternative to designing controllers by hand. Hand-design is difficult because it requires anticipating all possible operating conditions. Moreover, this laborious process has to be repeated whenever the configuration of the robot changes and it may even be impossible in some cases. In contrast, evolution can design a new controller automatically for the new configuration in simulation. For example, if a leg actuator fails on a quadruped robot during a remote mission, then it must continue the mission with minimum performance degradation by utilizing only the remaining three legs. ENSO evolved a straight and effective gait for such a three-legged configuration by disabling a leg in simulation. The resulting controller produced the same gait when transferred to the physical robot, suggesting that the approach presented in this paper can be utilized to design fault-tolerant robots for real-world applications.

The physical quadruped robot used in the experiments can be extended to a hexapod easily by adding two more legs to the attachments in the middle. This extension should be able to demonstrate that the ENSO approach scales up to a physical robot with more legs. ENSO can also be tested on robots with

more complex legs such as those with additional joints. Physically realistic simulations have shown that ENSO can evolve effective gaits, e.g. for a quadruped robot with knee joints [36]. The resulting well-coordinated movements that ENSO evolved suggest that it can be extended to evolve controllers for other multi-segmented robot morphologies as well, such as articulated arms and for serpentine robots. For these one-dimensional robots, the search space of symmetries can potentially be reduced by initializing ENSO with controller structures that represent their particular morphologies as closely as possible.

Another interesting extension is to attach Dynamixel AX-S1 sensors to the physical robot to measure distance [30]. Readings from these sensors can then be used as additional controller inputs, making it possible to evolve high-level behaviors that respond effectively to obstacles in the environment.

In addition to extending the robot, the ENSO approach itself can be extended in several ways to evolve controllers more effectively. First, the current manual decomposition of the controller into modules can be automated using hierarchical clustering algorithms based on the robot morphology. Second, instead of using a fixed neural network architecture for the modules, the architectures can be evolved using techniques such as NEAT [32]. Third, crossover of genotype trees can be implemented by swapping subtrees of parent trees if those subtrees have the same structure and node colors. Future work will focus on such extensions in order to eventually provide the sophistication necessary for evolving controllers for real-world robots.

6 Conclusion

This paper demonstrated how the ENSO approach can be utilized to design effective controllers for a physical quadruped robot. Controllers were evolved for an approximate model of the robot in a physical simulation using ENSO and the resulting controllers were then transferred to the physical robot. ENSO makes such transfer tractable by evolving symmetric neural network controllers. Since these controllers are actually coupled-cell systems, they produce stable gaits that are robust to inaccuracies in the simulation and uncertainties in the real world. ENSO evolved effective such controllers both for a fully functional version of the robot and for a version with a disabled leg. Moreover, the results of generalization experiments suggest that these controllers would be robust to common real-world challenges such as variations in battery performance and obstacles.

References

[1] R. D. Beer, H. J. Chiel, and L. S. Sterling. Heterogeneous neural networks for adaptive behavior in dynamic environments. In *Advances in Neural Information Processing Systems 1*, pages 577–585, 1989.

[2] R. D. Beer and J. C. Gallagher. Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.

[3] BigDog, 2010. http://www.bostondynamics.com/robot_bigdog.html.

[4] A. Billard and A. J. Ijspeert. Biologically inspired neural controllers for motor control in a quadruped robot. In *Proceedings of IJCNN*, pages 637–641, 2000.

[5] J. C. Bongard and H. Lipson. Once more unto the breach: Co-evolving a robot and its simulator. In *Proceedings of ALIFE9*, pages 57–62, 2004.

[6] R. A. Brooks. A robot that walks; emergent behaviors from a carefully evolved network. Technical Report AIM-1091, Massachusetts Institute of Technology, 1989.

[7] R. A. Brooks. Artificial life and real robots. In *Proceedings of the First European Conference on Artificial Life*, pages 3–10, 1992.

[8] J. E. Clark. *Design, Simulation, and Stability of a Hexapedal Running Robot*. PhD thesis, Department of Mechanical Engineering, Stanford University, 2004.

[9] J. J. Collins and I. N. Stewart. Coupled nonlinear oscillators and the symmetries of animal gaits. *Journal of Nonlinear Science*, 3(1):349–392, 1993.

[10] C. Ferrell. Failure recognition and fault tolerance of an autonomous robot. *Adaptive Behavior*, 2(4):375–398, 1994.

[11] D. Floreano and F. Mondada. Evolutionary neurocontrollers for autonomous mobile robots. *Neural Networks*, 11:1461–1478, 1998.

[12] L. Gao, S. Liu, and R. A. Dougal. Dynamic lithium-ion battery model for system simulation. *IEEE Transactions on Components and Packaging Technologies*, 25(3):495 – 505, September 2002.

[13] F. Gomez and R. Miikkulainen. Transfer of neuroevolved controllers in unstable domains. In *Proceedings of GECCO*, 2004.

[14] F. Gruau. Automatic definition of modular neural networks. *Adaptive Behavior*, 3(2):151–183, 1994.

[15] P. Holmes, R. J. Full, D. Koditschek, and J. Guckenheimer. The dynamics of legged locomotion: Models, analyses, and challenges. *SIAM Review*, 48(2):207–304, 2006.

[16] G. S. Hornby, S. Takamura, J. Yokono, O. Hanagata, T. Yamamoto, and M. Fujita. Evolving robust gaits with AIBO. In *Proceedings of ICRA*, volume 3, pages 3040–3045, 2000.

[17] N. Jakobi. *Minimal Simulations for Evolutionary Robotics*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, 1998.

[18] N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In F. Moran, A. Moreno, J. Merelo, and P. Chacon, editors, *Advances in Artificial Life*, volume 929 of *Lecture Notes in Computer Science*, pages 704–720, 1995.

[19] D. E. Koditschek, R. J. Full, and M. Buehler. Mechanical aspects of legged locomotion control. *Arthropod Structure and Development*, 33(3):251–272, July 2004.

[20] N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of ICRA*, volume 3, pages 2619–2624, May 2004.

[21] S. Koos, J.-B. Mouret, and S. Doncieux. Crossing the reality gap in evolutionary robotics by promoting transferable controllers. In *Proceedings of GECCO*, pages 119–126, 2010.

[22] H. Lipson, J. Bongard, V. Zykov, and E. Malone. Evolutionary robotics for legged machines: From simulation to physical reality. In *Proceedings of the 9th International Conference on Intelligent Autonomous Systems.*, pages 11–18, 2006.

[23] M. Mataric and D. Cliff. Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1):67–83, October 1996.

[24] O. Miglino, H. H. Lund, and S. Nolfi. Evolving mobile robots in simulated and real environments. *Artificial Life*, 2:417–434, 1995.

[25] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In *Proceedings of Artificial Life IV*, pages 190–197, 1994.

[26] Objet Eden 260V, 2010. `http://www.objet.com/3D-Printer/Eden260V/`.

[27] ODE: Open dynamics engine, 2007. `http://www.ode.org/`.

[28] OPAL: Open physics abstraction layer, 2007. `http://opal.sourceforge.net/`.

[29] M. H. Raibert. Legged robots. *Communications of the ACM*, 29(6):499–514, 1986.

[30] Robotis, 2010. `http://www.robotis.com/`.

[31] SolidWorks, 2010. `http://www.solidworks.com/`.

[32] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100, 2004.

[33] V. K. Valsalam. *Utilizing Symmetry in Evolutionary Design*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, 2010. Technical Report AI-10-04.

[34] V. K. Valsalam and R. Miikkulainen. Modular neuroevolution for multilegged locomotion. In *Proceedings of GECCO*, pages 265–272, 2008.

[35] V. K. Valsalam and R. Miikkulainen. Evolving symmetric and modular neural networks for distributed control. In *Proceedings of GECCO*, pages 731–738, 2009.

[36] V. K. Valsalam and R. Miikkulainen. Evolving symmetry for modular system design. *IEEE Transactions on Evolutionary Computation*, 15(3):368–386, 2011.

[37] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1 – 18, 2002.

[38] J. C. Zagal and J. Ruiz-Del-Solar. Combining simulation and reality in evolutionary robotics. *Journal of Intelligent and Robotic Systems*, 50(1):19–39, 2007.

[39] S. S. Zhang, K. Xu, and T. R. Jow. The low temperature performance of li-ion batteries. *Journal of Power Sources*, 115(1):137 – 140, 2003.

[40] V. Zykov, J. Bongard, and H. Lipson. Evolving dynamic gaits on a physical robot. In *Proceedings of GECCO, Late-Breaking Papers*, 2004.