

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**MORPHOLOGICAL DESIGN AND CONTROL OF A
BIO-INSPIRED,
STRUCTURALLY COMPLIANT QUADRUPED**

A thesis submitted in partial satisfaction of the
requirements for the degree of

Master of Science

in

COMPUTER ENGINEERING

by

Dawn Hustig-Schultz

June 2017

The thesis of Dawn Hustig-Schultz
is approved:

Professor Mircea Teodorescu, Chair

Professor Ricardo Sanfelice

Vytas SunSpiral

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures	v
List of Tables	vii
Abstract	viii
Dedication	x
Acknowledgments	xi
1 Introduction	1
2 Background	7
2.1 Rigid-Bodied Quadruped Robots	7
2.2 Compliant Quadruped Robots	7
2.3 Tensegrity Structures for Robotics	8
3 Methods	10
3.1 Mechanical Design	10
3.1.1 Reduction of Total Actuators	15
3.1.2 Integrated Spine and Leg Actuation	18
3.2 Machine Learning for Locomotion	18
3.2.1 Genetic Evolutionary Algorithms and Reinforcement Learning	19
3.3 Hierarchical Structure of Biological Nervous Systems	20
3.3.1 Impedance Controllers	20
3.3.2 Central Pattern Generators	21
3.4 Solution Space Reduction	23
4 Results and Discussion	24
4.1 Spine Driven Locomotion	24
4.1.1 Control	24
4.1.2 Mechanical Design	29

4.2	Reduced Actuation and Leg Control	29
4.2.1	Control	30
4.2.2	Solution Space	35
4.3	Integration of Spine and Leg Control	36
5	Conclusions and Future Work	41
	Bibliography	43

List of Figures

1.1	The Quadruped balancing on blocks, naturally adapting to complex footing by utilizing the multi-DOF compliance of its tensegrity spine.	2
1.2	Two different approaches to actuating MountainGoat.	5
3.1	The first two versions of MountainGoat. Machine learning was not applied to these two models, instead they were evaluated qualitatively, leading to improved models.	13
3.2	The second revision of MountainGoat, called LongTorso, with an extra vertebra added to the spine.	14
3.3	A close-up of the spine on Spirals, showing the extra cables added to the spine highlighted in red. Though these cables appear almost as duplicates of the previously existing cables, they only share one endpoint in common with each corresponding cable from the original spine morphology. . . .	15
3.4	Model of MountainGoat, <i>NoFeet</i> with feet removed.	16
3.5	Two different actuator configurations consisting of spiral muscles, which were first mentioned in section 3.1. Both of these actuator configurations produce torsion in the spine, which assists in lifting the legs of the robot. Since there are more active cables in the configuration in figure 3.5a, it produces more torsion than the configuration in figure 3.5b. Consequently, it will be shown in section 4.2.1 that the first configuration produces faster locomotion speed.	17
4.1	A comparison of the results from the Monte Carlo stage of machine learning for LongTorso and Spirals.	25
4.2	The initial 30,000 randomized Monte Carlo trials, each one minute in duration, for NoFeet. For the evolutionary results of the best 40 of these trials, see table 4.2.	26
4.3	The evolution process and the resulting trajectory of MountainGoat. . .	27
4.4	Static prototype of NoFeet.	29
4.5	A comparison of the Monte Carlo results for SpiralsOnly and Reduced-Spirals.	31

4.6	The initial 90,000 randomized Monte Carlo trials, each one minute in duration, for Achilles. For the results of the best 40 of these trials, see table 4.2 and figure 4.7a.	33
4.7	The evolution process and the resulting trajectory of MountainGoat. . .	34
4.8	The initial 90,000 randomized Monte Carlo trials, each one minute in duration, for the model using both active spine and leg muscles.	37
4.9	Evolution and Trajectory of SpiralsOnly, Achilles, and AchillesSpirals, to show improvement when both are combined.	39

List of Tables

3.1	Parameters used for each model	19
4.1	Longest Distance Per Minute for Each Model	26
4.2	Longest Distance Per Minute for Each Model	34
4.3	Number of parameter pairs learned for each model	36
4.4	Longest Distance for Achilles and Number of Parameter Pairs	38

Abstract

Morphological Design and Control of a Bio-Inspired,
Structurally Compliant Quadruped

by

Dawn Hustig-Schultz

From the viewpoint of evolution, vertebrates first accomplished locomotion via motion of the spine. Legs evolved later, to enhance mobility, but the spine remains central. Contrary to this, most robots have rigid torsos and rely primarily on movement of the legs for mobility. The force distributing properties of tensegrity structures presents a potential means of developing compliant spines for legged robots, with the goal of driving motion from the robots core. In addition, the increasing complexity of soft and hybrid-soft robots highlights the need for more efficient methods of minimizing machine learning solution spaces, and creative ways to ease the process of rapid prototyping.

In this thesis I present the process of morphological design for a tensegrity quadruped robot, the first to the author's knowledge, which I call MountainGoat, and its impact on controllable locomotion. All parts of the robot, including legs and spine, are compliant. Control is initially demonstrated with three variations of MountainGoat, focusing on actuation of the spine as central to the locomotion process. Following the general pattern of biological evolution, leg actuation is developed next. Additionally, to reduce the overall machine learning space, I present four different choices of muscle groups to actuate: three for a primarily spine-driven morphology of a tensegrity

quadruped, called MountainGoat, and one for a primarily leg-driven variation of this quadruped, and compare the resulting differences in locomotion speed. Each iteration of design seeks to reduce the total number of active muscles, and consequently reduce the dimensionality of the problem for machine learning, while still producing effective locomotion. The reduction in active muscles seeks to simplify future rapid prototyping of the robot. For this portion of the thesis, two separate approaches to actuation, one primarily spine-driven and the other primarily leg-driven, are explored.

Locomotion for all models is aided by the use of central pattern generators, feedback control via a neural network, and a two-tiered machine learning approach involving the Monte Carlo method as well as genetic evolution for parameter optimization.

This work is dedicated to my Father, Charles H. Hustig, who was my earliest engineering influence, and to my husband Kevin Schultz, whose love and support has helped sustain me.

Acknowledgments

I would like to thank Vytas SunSpiral, Brian Mirletz, and Perry Bhandal for developing and advising on the NTRT learning and machine learning framework, As well as others affiliated with the NASA Ames Intelligent Robotics Group, including Adrian Agogino, Andrew Sabelhaus, and Jonathan Bruce for further insights and support. Additional thanks goes to Tom Flemons for his initial quadruped design and inspiration.

I would also like to thank Paloma Fautley, Jian Hao Miao, Kevin Le, and Joshua Gier for their work on schematics, cad, and construction of the static prototype, and other members of DANSER Lab at UCSC, including Nick Cramer, Steve Lessard, Dennis Castro, Amir Pourshafiee, Sina Kahnemouyi, Samira Zare, Anca Popescu, Gordon Keller, and Calvin Chopra for their general help, encouragement, and support. And I would expressly like to thank professor Mircea Teodorescu, my advisor for providing 3 years of guidance and insight in my research career.

The text of this thesis includes reprints of the following previously published material: [[14]]

Chapter 1

Introduction

Current wheeled and rigid-bodied robots are limited in their mobility over rough terrains and in their ability to operate in unpredictable environments. This limits their usefulness for such tasks as space exploration, search and rescue missions, and missions in environments unsafe for humans. These environments are well suited for compliant quadruped robots.

Class-1 Tensegrity structures consist of disjoint compression members (rods) interconnected by a system of tension members (cables) with no two compression members in direct contact. Other class-k structures, with k equaling the number of compressive members that come into contact at a movable joint, also exist [38]. These structures were introduced in architecture by Kenneth Snelson [39], but more recently have been incorporated into robotics. Tensegrity robots are lightweight and robust to failures, as the failure of one actuator leads to diminished performance rather than failure of performance. They are impact tolerant, as forces distribute evenly over the whole instead

of being magnified into joints by internal lever arms, causing less damage to itself and to other objects in its environment.

Figure 1.1 shows a model of MountainGoat, a tensegrity quadruped robot based off an original model by Tom Flemons [9], balanced on a terrain filled with blocks. Of note in this figure is MountainGoat’s passive terrain interaction, and how it naturally adapts to complex footing by utilizing the multiple degree-of-freedom compliance of its tensegrity spine. This ability of tensegrities to redistribute forces to achieve equilibrium is a compelling reason for their application to constructing robots that can traverse rugged ground. The goal of this research is to develop a quadruped robot with the agility and adaptability of a mountain goat.




Figure 1.1: The Quadruped balancing on blocks, naturally adapting to complex footing by utilizing the multi-DOF compliance of its tensegrity spine.

Prior studies with other tensegrity morphologies has shown robust locomotion


over rough terrain [42], [15], [2], [25], [22], [23], [24], [21], [40]. The quadruped morphologies we present, however, have not yet been developed to the point where they can locomote over rugged ground. With the results presented here, we are beginning to understand the process of whole-body control, and how the spine of a compliant robot provides support to shoulders and hips in order to lift legs. For instance, we have not quite achieved the amount of leg lift necessary to actively traverse rugged terrain. In addition, our quadruped still lacks knees, which enhance motion over obstacles. Nevertheless, these findings have given important insights toward the ultimate goal of a compliant quadruped that can travel over multiple types of terrains.

As tensegrity structures and their control mechanisms become more complex, it is important to find strategies for keeping the dimensionality of the solution space small for more efficient performance of and more successful outcomes from machine learning. Iscen, et. al. have explored the use of coevolutionary algorithms to control underactuated tensegrity structures [16], while keeping solution space small. Previous research into the redundancy of actuators on the SUPERball tensegrity robot, by Lessard, et. al. has also shown that reduction in active muscles gracefully degraded productive locomotion, rather than curtailing it [17]. This study showed that the reduction of the dimensionality of the solution space for a tensegrity robot, which can help increase the chances of finding desirable solutions, will not necessarily harm the robot's performance and can ease the process of learning desirable locomotion solutions.


Because of the increasing complexity of tensegrity robots, rapid prototyping can be more challenging than for rigid-bodied robots. One major obstacle is the large

amount of actuated cables that may be needed in order for a tensegrity quadruped to move effectively even over flat terrain. These actuated cables sinusoidally change length in order to make the robot move. On a robotic prototype, the mechanism for this changing of length could be the spooling of cable by a DC motor. This effect could also be achieved using pneumatic actuators, Shape Memory Alloys (SMAs), or dielectric elastomers. Some promising potential solutions exist in the form of pneumatic actuators, as investigated by Polygerinos, et. al. [32] and Niiyama, et. al. [27], Shape Memory Alloys as implemented by Umedachi, et. al. [41], and dielectric elastomers, as explored by Pelrine, et. al. [30], Petralia et. al. [31], and Bilodeau, et. al. [4]. Many of these solutions, however, have drawbacks, such as heat dissipation and reliability for SMAs or the need for large volumes of air for pneumatics. Others, such as soft actuators, are still in early developmental stages. Initial rapid prototyping could potentially benefit from being able to use readily available, off-the-shelf parts, with reliance on more sophisticated actuators left to later prototypes. For more complex robotic morphologies, this means finding the most efficient approach that minimizes the use of commercially available actuators.

For our exploration of minimized solution space dimensionality, we will use two different quadruped morphologies. In our previous publication, we have discussed the centrality of the spine to locomotion [14], using the morphology seen in figure 1.2a. Besides the spine's central role in locomotion, a well-formed achilles tendon also plays an important supporting role. Various research groups have found that, though the gastrocnemius muscle does little mechanical work, elastic energy storage and return from



(a) This version of MountainGoat will be used for spine actuation only, implementing three different actuator configurations that use 52, 24, and 16 muscles. See figures 4.2 4.5a and 4.5b for simulation results, and figure 4.7b for the different locomotion trajectories.



(b) This version of MountainGoat, called Achilles, will use leg actuation only. The muscles to be actuated are highlighted on only two of the legs in the figure, but will be actuated on all four legs. See figures 4.6 for simulation results and 4.7b for the locomotion trajectory.

Figure 1.2: Two different approaches to actuating MountainGoat.

the achilles tendon plays a major role in power production at the ankle of bipeds and quadrupeds [8], and that the absence of a well developed achilles tendon muscle would prevent bipeds from running effectively both at high speeds and over long distances [37]. This suggests the importance of having an achilles tendon to help produce the necessary ground reaction force for quadruped locomotion. Researchers at MIT have explored the advantages of an achilles tendon in the Cheetah robot [3]. In the process of reducing the solution space of our robot, we will also present preliminary findings of an implementation of an achilles tendon on MountainGoat, for which we will simulate the morphology shown in figure 1.2b. We will focus on a total of four different hand-chosen actuation solutions for the two morphologies mentioned above, one for the leg-driven morphology, and three for the spine-driven morphology.

Chapter 2

Background

2.1 Rigid-Bodied Quadruped Robots

Boston Dynamics' BigDog and Spot robots have had success in navigating robust terrain including on ice and snow [33] [1]. These robots can be energy expensive, prone to single-point failures, and susceptible to damage on impact, to the robot itself as well as to objects and people in its environment. Degraeve, et. al. have incorporated some compliance in the legs of quadruped robots [6]. These rigid bodied robots, however, represent more constrained solutions that lack the compliant spines that are central to the speed, agility, and stability of quadruped and biped locomotion [12].

2.2 Compliant Quadruped Robots

The benefit of a compliant spine to quadruped locomotion has been studied by Zhao, et. al., simulating robots that have multiple spinal joints ranging in number

from 1 to 12 [43]. Although improved locomotion was shown with two and four spinal joints, these simulated designs use one-DOF joints that represent single points of failure, and the legs of the robots in these simulations were completely rigid. Researchers at the University of Pennsylvania compared two robots with the same semi-rigid c-shaped legs, one with a rigid body and the other with a parallel elastic actuated spine. The robot with the elastic spine showed more distance and agility in forward leaps than its rigid bodied counterpart [7].

2.3 Tensegrity Structures for Robotics

One of the earliest investigations of tensegrity locomotion involved gait production in a simple three-bar tensegrity structure by researchers at Cornell University [29]. Some robots, such as MIT's Cheetah robot, have incorporated tensegrity principles in the legs, but not in the spine of the robot, where it could have greater benefit [36]. Although the legs of the Cheetah robot are very effective in forward motion, they are somewhat limited in the kind of lateral motion needed to give good balance and stability.

Various morphologies can incorporate tensegrity structures. Xydes, et. al. have studied the locomotion of a snake-like tensegrity structure for duct inspection [42], Lessard, et. al. have developed tensegrity arms [18], and Agogino, et. al. have studied locomotion of the SUPERBall robot, a tensegrity rover intended for exploration of Saturn's moon Titan [15], [2]. Tensegrity structures can be used to model spines[19],

and extensive research has been done by Mirlatz, et. al. on flexible tensegrity spines for their potential benefit in locomotion, with the eventual goal of building compliant quadruped and biped robots [25], [22], [23], [24], [21], [40]. The study of tensegrity spines in these papers has demonstrated the robustness of tensegrity locomotion over rough terrain. Similar spines have been incorporated into rigid-legged quadruped robots by Sabelhaus, et. al. [35].

Chapter 3

Methods

3.1 Mechanical Design

All our models are class-1 tensegrity structures, since rods are only connected to each other via cables. Our structural design approach began with a fully passive model of MountainGoat, designed by Tom Flemons [9]. Due to the tendency of tensegrity structures to redistribute loads and deform to equilibrium shapes, designing for movement can be very counter intuitive, and structural design and control end up being highly coupled. Because of this property, many of our design iterations came about as a result of attempts to control previous model designs.


The initial design, called Flemons, seen in figure 3.1a, consists of a spine of six X-segments of four struts each, with three of the segments oriented vertically and three oriented horizontally. ten cables, which play a similar role to muscles in biological creatures and which serve as the actuators, connect each segment in the spine, with the

exception of the penultimate segment which is connected by fourteen cables. The hips and shoulders consist of single T-segments, of three struts each. Each of the shoulders are connected to the spine by ten cables, and each of the hips are attached to the spine by nine cables. the hips connect to each other with one cable, to add stability. Legs also consisted of single X-segments, but with shorter support rods added to the bottom of each leg for stability. Each leg is connected to its corresponding hip or shoulder, foot, and spine by fourteen cables. The feet consist of two rods which cross each other, and 11 cables connecting these rods to the legs, for support. The full model has a total of 60 struts. In qualitative testing on hilly terrain, we found that the cross rods of the foot caught too easily on obstacles, and didn't provide enough stability to keep the structure standing.


Figure 3.1b shows NewFeet, which differs only in the feet. The bottom rod of the leg is extended by 5 cm, and a compliant foot, consisting of a four strut prism, replaces the two cross rods, and brings the model to a total of 68 struts. An extra cable was also added to each hind leg, shown in red in figure 3.1b, connecting these legs to the last vertebra in the spine to improve overall balance. This compliant foot provided more stability to the structure. To test this new model, we used a passive drop test. The drop test was done in simulation by placing the robot at a starting position so that the feet are 30 cm higher than the ground. For context, the robot itself is 36 cm tall. Two different terrains were employed. The first consists of randomly placed 5 cm wide by 5 cm tall by 5 cm long blocks, and the second consisted of evenly spaced 6 cm tall hills with a base diameter of 6 cm. The simulation is then run in graphics mode,

so that the final landing position can be observed. This was then repeated multiple times. These tests are relevant, as they allowed us to assess the passive stability of the robot's morphology, which helps assure that the robot will also be stable during actuation. In drop tests on hilly and block-filled terrains, these feet helped the model to maintain a standing position on landing, rarely falling over. Flemons and NewFeet were not actuated, but were primarily evaluated passively and qualitatively for structural stability.

For ease of actuation, the structure was updated again by adding an extra vertebra to the spine, bringing the total number of struts to 72. This was done to simplify the initial approach to control, as each vertebra of the spine would have an equal number of CPGs. Then input parameters for only 8 CPGs on one vertebra would need to be learned, and then applied to the 7 identical spine segments, reducing the overall solution space. In comparison, Flemons and NewFeet would need to learn parameters for an additional 12 cables, to accommodate the penultimate vertebra. The change had the extra benefit of providing enough distance between the front and back feet to keep them from colliding with each other as well as adding a bit more torsion to the spine. Torsion is the twisting of the spine of a quadruped in the transverse plane, that is, the plane that divides the quadruped into anterior and posterior portions. For our simulations in NTRT, this consists of the xz-plane. This motion is important because it helps to pull the legs vertically, off the ground, in a regular alternating pattern. Figure 3.2 shows this new model of MountainGoat, called LongTorso. LongTorso has a total of 56 CPGs, in the spine only, which actuate the simulated model. Results of spine



(a) The initial model of the quadruped, Flemons, based off a design by Tom Flemons. On uneven ground, the feet didn't provide the desired structural support and were prone to catching on obstacles.



(b) The first revision of MountainGoat, called NewFeet, with improved compliant feet for better balance and reduced catching on obstacles. Differences are shown in red.

Figure 3.1: The first two versions of MountainGoat. Machine learning was not applied to these two models, instead they were evaluated qualitatively, leading to improved models.

actuation on this model can be seen in section 4.1.1.




Figure 3.2: The second revision of MountainGoat, called LongTorso, with an extra vertebra added to the spine.

As stated earlier, morphological design and actuation of a tensegrity robot are highly coupled, due to the tendency of these structures to redistribute loads and deform to equilibrium shapes. As an example, early attempts at manually designing a controller to lift a single leg resulted in the corresponding shoulder drooping toward the foot, rather than the leg lifting off the ground. Since the initial cable layout of the spine did not provide the stability to hold the shoulder up, two extra spirals of cables were added to the spine, one clockwise and one counterclockwise. These extra twenty cables, which serve a similar function as the latissimus dorsi muscles in many four-legged vertebrates, can be seen on Spirals in figure 3.3. The addition of these cables helped increase torsion in the spine, which led to increased distance traveled in simulation, as

is shown in section 4.1.1.




Figure 3.3: A close-up of the spine on Spirals, showing the extra cables added to the spine highlighted in red. Though these cables appear almost as duplicates of the previously existing cables, they only share one endpoint in common with each corresponding cable from the original spine morphology.

Simulations of control of the structures in figure 3.2 led to the conclusion that ground reaction force was being lost in the compliance of the feet. As shown on *NoFeet* in figure 3.4, the feet were thus removed, and two extra cables were added between each leg and its adjacent body segment, to keep the model standing. This removal, which reduced the total amount of struts to 56, added more distance to simulated locomotion experiments, as shown in section 4.1.1. LongTorso, Spirals, and NoFeet are 86 cm long by 42 cm wide by 36 cm tall. Flemons and NewFeet have a similar scale.

3.1.1 Reduction of Total Actuators

Our first actuator configuration of NoFeet, which we discussed in section 3.1 consists of all 52 spine cables as shown in figure 1.2a. This design will be included as a comparison to other configurations with fewer actuators.

Figure 3.5a shows a configuration, which we will henceforth refer to as Spiral-





Figure 3.4: Model of MountainGoat, *NoFeet* with feet removed.

Only, consisting of 24 actuators, shown in red. The passive cables, left unhighlighted, were left in the morphology in order to provide support to the structure. This passive function is similar to that of fascia in biological creatures. The choice to actuate only these cables followed naturally from this beneficial function, which in our previous work gave extra torsion and support to the shoulders of MountainGoat. The question of interest in making this choice is whether actuating this set of spiral cables alone will degrade performance, as previous studies have shown [17], while still allowing the morphology to still successfully carry out its intended function, or whether some improvement will be gained by reducing the interference of overlapping actuators.


The third actuator configuration that will be explored can be seen in figure 3.5b. This configuration, which will be called ReducedSpirals, consists of 16 actuators in green. It is essentially the same as the configuration of 24 actuators from figure 3.5a, but with the first and last four actuators converted to passive fascia. Again, the

question of interest is whether the reduction in the number of active spiral cables will gracefully or severely degrade performance.

Figure 1.2b shows an actuator configuration that depends on leg cables, rather than on cables in the spine. This configuration, called Achilles, has an opposing pair of cables in each leg, which act similarly to the combination of the gastrocnemius muscle and the Achilles tendon in many quadrupeds and bipeds. Since there are only two active leg cables, the total number of actuators for this configuration is 8. Additional rods in the feet of the model were needed to increase stability during locomotion.



(a) An illustration of the spirals, which consist of 24 active cables, shown in red.



(b) An illustration of the shorter spirals, which consist of 16 active cables, shown in green.

Figure 3.5: Two different actuator configurations consisting of spiral muscles, which were first mentioned in section 3.1. Both of these actuator configurations produce torsion in the spine, which assists in lifting the legs of the robot. Since there are more active cables in the configuration in figure 3.5a, it produces more torsion than the configuration in figure 3.5b. Consequently, it will be shown in section 4.2.1 that the first configuration produces faster locomotion speed.

3.1.2 Integrated Spine and Leg Actuation

Considering the centrality of the spine for locomotion, as hypothesized by Gracovetsky, et. al. [12], as well as the importance of the Achilles tendon for push off, as extensively explored by Folkertsma, et. al. [10], we wanted to see how combining active spine and leg cables would affect locomotion speed. SpiralsOnly was chosen as the spine actuator configuration to combine with leg actuation since, as will be shown in section 4.2.1, it was the fastest of the three spine configurations. We will refer to this combined leg and spine actuator system as AchillesSpirals.

3.2 Machine Learning for Locomotion

The open source NASA Tensegrity Robotics Toolkit (NTRT) was used for simulation ¹. NTRT is built on the Bullet Physics Engine, version 2.82, which handles rigid body dynamics to simulate the rods of the structure. This is supplemented by an additional custom soft body spring-cable model with contact dynamics, which is used to simulate the cables of the structure. The dynamics of the spring-cable are based on Hooke’s law for a linear spring, and collisions are detected using ghost objects within Bullet [24]. Internal cable and rigid body dynamics were previously validated within 1.3% error [5]. Additional tests validated steady state error on maximum cable tension within 6.1%, maximum system tension on hand-tuned controllers was validated within 7.9% error, and tensions from CPGs were validated within 1.6% error [24].

¹Source code for NTRT can be found at <https://github.com/NASA-Tensegrity-Robotics-Toolkit>

Simulations were run at 1000 Hz. Five different iterations of MountainGoat are presented. The stiffness, pretension, and damping parameters used for these models can be found in table 3.1. For reference, a pretension setting of 700 is about 5N, 1000 is about 7N, 2500 is about 17.9N, 3500 is about 25N and 10,000 is about 71N. Damping is in kg/cm^3 and Stiffness is in kg/s^2 . The last three of these model revisions were tested in simulation on flat terrain, using the actuation approach discussed below.

Table 3.1: Parameters used for each model

Model	Parts	pretension	stiffness	damping
Model 1: Flemons	All parts	700	2000	20
Model 2: NewFeet	All parts	700	2000	20
Model 3: Long- Torso	Spine	0	1000	10
	Legs	2500	3000	30
	Feet	1000	1000	10
Model 4: Spirals	Spine	0	1000	10
	Legs	3500	4000	10
	Feet	1000	4000	10
Model 5: NoFeet	Spine	0	2000	10
	Legs	4000	10000	10

3.2.1 Genetic Evolutionary Algorithms and Reinforcement Learning

Inspired by evolution, and the central role spines play in vertebrate locomotion, we initially focused on driving motion from the spine. This approach allowed us to continue the research of Brian Mirletz [25], [22], [23], [24], [21], [40], by applying his tensegrity spine control research to MountainGoat. Thus, we use machine learning to optimize the controls for novel morphologies, allowing us to evaluate the effectiveness

of a specific morphology. These methods were also adapted to control the legs of the robot, in order to explore the role of the Achilles tendon in locomotion.

A learning run starts with a Monte Carlo stage, where 30,000 random trials are generated. Each trial has a duration of 60 seconds and, since we use the distance traveled in one minute as the measure of fitness, this distance is determined by taking the difference between the location of the center of mass at the beginning and end of each trial. Those trials in which MountainGoat travels the greatest distance in any direction are considered the fittest trials. After the Monte Carlo stage ends, we evolve the fittest 40 trials via a genetic algorithm with crossover, mutation, and elitism. Given a complex parameter space with many peaks and valleys, we use Monte Carlo trials to find initial decent results, and then optimize them with the genetic algorithm. The mutation chance used was 50% while the mutation deviation was 3%.

3.3 Hierarchical Structure of Biological Nervous Systems

Our approach to actuation reflects the hierarchical nature of biological nervous systems, with local reflexes at a lower level and Central Pattern Generators (CPGs) at a higher level [13].

3.3.1 Impedance Controllers

Impedance control is used for the lower level reflexes, based on an equation first used for tensegrity by Orki, et. al. [28] and adapted to account for descending commands from the CPGs by Mirletz, et. al. [21]:

$$T = T_0 + K(L - L_0) + B(V - V_0) \quad (1)$$

T is the output tension, T_0 is the tension offset, and K is the position gain on the difference between the current length L and the desired length L_0 . B is the velocity gain on the difference between the current velocity V and the desired trajectory V_0 . This V_0 term is a descending command from the CPG.

3.3.2 Central Pattern Generators

The CPG equations used consist of adaptive phase coupled oscillator equations with frequency feedback [34], as well as amplitude and phase feedback [11], and were previously used by Mirlatz, et. al. for locomotion of tensegrity spines [25]:

$$\dot{r}_i = \gamma(R_i + k_r F_r - r_i^2)r_i \quad (2.1)$$

$$\dot{\theta}_i = \omega_i + k_\theta F_\theta + \sum_j r_j w_{ij} \sin(\theta_j - \theta_i - \phi_{ij}) \quad (2.2)$$

$$\dot{\omega}_i = k_\omega F_\omega \sin(\theta_i) \quad (2.3)$$

$$\dot{V}_i = r_i \cos(\theta_i) \quad (2.4)$$

where r_i is the wave's amplitude, ω_i is its frequency, and θ_i is its phase. V_i is the input to the impedance controller. The amplitude, as seen in equation 2.1, is set by the convergence parameter γ and the setpoint R_i . The coupling weight w_{ij} , phase

offset ϕ_{ij} , and the the amplitude of the neighboring node r_j are used to determine the derivative of the phase $\dot{\theta}_i$. The resulting V_i from equation 2.4 is used as input to the lower level impedance controller. The constant terms k_r , k_θ , and k_ω are scalar gains on the corresponding feedback parameters F_r , F_θ , and F_ω . These feedback parameters come from outputs of an artificial neural network, and are used in a similar manner as in [11] and [25]. As in [25], this neural network consists of two input nodes for which the inputs are tension and length, one hidden layer of four nodes, and three output nodes for F_r , F_θ , and F_ω .

To keep the size of the solution space manageable, the two parameters explored by the algorithm for all models are the phase offset ϕ_{ij} and coupling weight w_{ij} , from equation (2.2). Our leg actuator configuration, Achilles, which employs antagonistic pairs of long cables, required re-tuning of the PD and impedance controller parameters T_0 , K , and B from equation (1), which previously were empirically tuned for shorter cables. While these parameters can be tuned by hand, requiring time, trial, and error, our machine learning algorithm will instead be used to tune these parameters for the Achilles model. Mirletz, et. al. experimented previously with machine-learned Impedance control parameters for tension, length, and velocity [25]. In addition to this, the proportion and derivative input parameters for PID control are learned via evolution of randomized trials.

The CPG equations are integrated using ODEInt, which is part of the Boost C++ libraries [26]. Each actuator is coupled only to other actuators that share rigid bodies [25], [21]. The linearity of the rigid bodies attached to these actuators ensures

that there are at most three rigid bodies in each coupling set.

3.4 Solution Space Reduction

For the Achilles model, we used a similar approach for reducing the solution space as in [21]. That is, our solution space consists of the coupling weights and phase offsets of one leg, and those parameters were reused for all four legs. This approach could be modified to produce different gaits by adding an offset pair for each leg to the individual weight-coupling pairs. The result would be the generation of different gaits, while still keeping the solution space small. Currently, our research does not yet utilize such a gait-dependent offset pair, but future work could explore different gaits using this technique.

Chapter 4


Results and Discussion

4.1 Spine Driven Locomotion

4.1.1 Control

Due to the counter intuitive nature of how these structures move, we found that hand-designing controllers was ineffective, so we turned to machine learning techniques, as originally developed by Mirletz for learning the control for tensegrity spines. Machine learning using CPGs and a neural network for feedback, as discussed in section 3.2, was used to actuate LongTorso, Spirals, and NoFeet from figures 3.2, 3.3, and 3.4 in simulation. Figure 4.1a shows the results for LongTorso from figure 3.2. As can be seen, only five of the 30,000 controllers achieved a distance greater than 30 cm/min, and only one of these controllers yielded a distance greater than 50 cm/min. Table 4.2 shows that genetic evolution more than doubled this distance, to 110.9 cm/min.

Figure 4.1b shows the results of 30,000 Monte Carlo trials for Spirals, from



(a) Distance traveled by LongTorso over 30000 Monte Carlo trials. (b) Distance traveled by Spirals over 30000 Monte Carlo trials.

Figure 4.1: A comparison of the results from the Monte Carlo stage of machine learning for LongTorso and Spirals.

figure 3.3. Although the extra spirals of cables added to Spirals did not improve the longest distance traveled over 30,000 Monte Carlo trials, the total number of Monte Carlo trials that traveled farther than 30 cm/min increased from 5 to 18, and the number of trials that achieved close to 50 cm/min increased from 1 to 4. Table 4.2 shows that evolution yielded a more than 100% improvement with a distance of 130.4 cm/min. This shows that this change in morphology, with the aim of increasing torsion in the spine, improved distance outcomes in CPG control.

NoFeet, from figure 4.2, showed the best results over both Monte Carlo and evolution. The longest distance yielded from Monte Carlo, which was 247.1 cm/min, was more than double the longest distance of both LongTorso and Spirals after 30,000 Monte Carlo trials. As can be seen in figure 4.2, ten of the trials traveled a distance greater than 100 cm/min, and many more traveled distances greater than 50 cm/min than for LongTorso and for Spirals from figures 3.2 and 3.3. Table 4.1 shows that

genetic evolution yielded a distance of 339.4 cm/min which is, again, more than double the distance of Spirals and more than triple the distance of LongTorso after genetic evolution. These gains from removing the feet show that distance was indeed lost from having feet with too much compliance. The resulting increase in ground reaction force from removing the feet led to increased speed.





Figure 4.2: The initial 30,000 randomized Monte Carlo trials, each one minute in duration, for NoFeet. For the evolutionary results of the best 40 of these trials, see table 4.2.

Table 4.1: Longest Distance Per Minute for Each Model

Model	Monte Carlo	Evolution
LongTorso	53.7 cm/min	110.9 cm/min
Spirals	51.4 cm/min	130.4 cm/min
NoFeet	247.1 cm/min	339.4 cm/min

Figure 4.3a shows the increase in distance over subsequent generations of evolution. This particular plot comes from the evolution of the best 40 trials from LongTorso,



(a) The distance traveled over 60 generations using control policies dictated by the 40 most successful Monte Carlo trials for LongTorso. (b) Trajectory of NoFeet's center of mass over a one-minute simulation. Note the straight path taken by the robot during locomotion.

Figure 4.3: The evolution process and the resulting trajectory of MountainGoat.

shown in figure 3.2, and for which the Monte Carlo results are shown in figure 4.1a. Note that the total distance increases more quickly over the first 20 generations than it does over the last 40 generations. This shows an example of how the machine learning and control scheme described in section 3.2 converges to a near-optimal locomotion gait for a given morphology.

Figure 4.3b shows the trajectory of the center of mass of NoFeet, from figure 3.4, over one minute. This shows that the simulated robot takes a straight path during travel. While the robot does not yet move quickly, we are still developing useful theories of morphological design. The current model of MountainGoat does not yet even have knees, for instance, which would enhance the robot's ability to move over uneven terrain. We have, however, shown with our above results the importance of spinal torsion for locomotion.

The machine learning results shown in figures 4.1a, 4.1b, and 4.2 and in table

4.1 yield interesting results about the complexities of morphological design of tensegrity structures for locomotion. The passive compliance of the tensegrity quadruped is very valuable, as it allows for natural force distribution and passive terrain adaptation. Yet passive compliance in some cases hinders the effectiveness of locomotion, as can be noted from the differences in the Monte Carlo results for Spirals and NoFeet, in figures 4.1b, and 4.2, as well as the longest distances traveled in table 4.1. More productive motion was gained from reducing compliance in this part of MountainGoat. A lesson learned from this outcome is that when it comes to leg design, especially when it comes to the future work of designing knees, would be that the legs should become mechanically stiff (the knees and hips "lock") during the stride phase and should become passively compliant (knees and hips "unlock") during the swing phase.

The advantage of tensegrity robots is that pretension and stiffness can have different settings in various body parts to enable more productive motion, as shown in table 3.1, where the legs of LongTorso, Spirals, and NoFeet have greater stiffness and pretension settings than the spine does. The ability for components of the robot to be either compliant or stiff is a unique characteristic of tensegrity robots. But, the morphology of the structure must be designed correctly to be able to provide those points of stiffness, such as lifting the shoulder via the spiral spine cables added to Spirals, from figure 3.3, which increased the distance traveled after Monte Carlo and genetic evolution. How to design for both passive compliance and active stiffness is an open research topic, for which there is currently no guiding theory.

4.1.2 Mechanical Design

To determine whether NoFeet from figure 3.4 represents a stable structure outside of simulation, we constructed a static prototype. Figure 4.4 shows this prototype, which consists of wooden dowels, plastic golf balls, and elastic strings held in place with brass hooks. Passive equilibrium due to force distribution is key in tensegrity structures, and this static prototype shows that NoFeet stands in a stable position the same way it does in simulation.




Figure 4.4: Static prototype of NoFeet.

4.2 Reduced Actuation and Leg Control


The three new actuator configurations discussed in this section represent ways of reducing the total number of actuators required for productive motion. This includes the two spine-only configurations, which include SpiralsOnly (24 active cables) and

ReducedSpirals (16 active cables), and the leg-only configuration – Achilles – which has only 8 active cables in total. The results of NoFeet presented in section 4.1.1, which has 52 active spine muscles, are included in the discussion as a reference point for comparing the performances of all the configurations that employ reduced actuation.

4.2.1 Control

The results of 30,000 Monte Carlo trials on the actuator configuration of NoFeet, illustrated in figure 1.2a above, can be seen in figure 4.2, which was previously discussed in section 4.1.1. This configuration of 52 active cables, yielded longer distances after Monte Carlo than SpiralsOnly or ReducedSpirals, with 24 actuators and 16 actuators, respectively. NoFeet produced many more trials which traveled farther than the 50 cm/minute threshold than for SpiralsOnly or ReducedSpirals. The distance traveled after evolution, which as we discussed above was 339.4 cm/min, is listed again for comparison in table 4.2.

The Monte Carlo results for SpiralsOnly, the configuration with 24 active spine cables as depicted in figure 3.5a, can be seen in figure 4.5a. The number of trials that traveled longer than 50 cm/min is much sparser than for NoFeet, which is in keeping with the idea explored by Lessard, et. al. [17]. However, the best trial traveled at a speed of 229.4 cm/min, which though slower than the best trial for NoFeet, as seen in figure 4.2, is nonetheless close. Despite the sparsity of trials faster than 50 cm/min, this configuration of 24 active spine cables improved the most during evolution. This evolution result could be due to the fact that SpiralsOnly had at least one more trial



(a) The initial 30,000 randomized Monte Carlo trials, each one minute in duration, for SpiralsOnly. See table 4.2 and figure 4.7a for the resulting distance after evolution from the best 40 of these trials.

(b) The initial 30,000 randomized Monte Carlo trials, each one minute in duration, for ReducedSpirals. See table 4.2 and figure 4.7a for the resulting distance after evolution from the best 40 of these trials.

Figure 4.5: A comparison of the Monte Carlo results for SpiralsOnly and ReducedSpirals.

with a distance longer than 150 cm/minute than NoFeet had. The favorable traits from this extra MC trial would most likely give it an advantage over multiple generations. As seen in table 4.2 and figure 4.7a, the distance traveled increased to 366.8 cm/min after the evolution phase of learning.

Figure 4.5b shows the results of 30,000 Monte Carlo trials for ReducedSpirals, the configuration of 16 active spine cables shown in figure 3.5b above. As with the configuration of 24 spine actuators, there is a sparser number of trials faster than the 50 cm/min threshold than there is for NoFeet. The use of 16 spine actuators resulted in a best trial of 150.2 cm/minute, which was slower than the best trials for both SpiralsOnly and NoFeet. The results of evolution continued this trend, increasing the speed of locomotion to 274.7 cm/min as in table 4.2 and figure 4.7a. In our previous work, we discussed the importance of this type of spiral cable configuration for increasing

torsion in the spine and giving more shoulder support for lifting the legs [14]. The resulting performance of ReducedSpirals, in comparison to SpiralsOnly, supports this claim, inasmuch as decreasing the length of this spiral lead to decreased torsion and shoulder support. Also, the general reduction in the number of actuators supports the graceful degradation, rather than sharp cessation, in performance.

Since we use machine learning to tune the PD and impedance parameters for Achilles, we ran three times the amount of Monte Carlo trials. Recall from section 3.1.1 that Achilles moves by means of actuating an opposing pair of cables in each leg, rather than by means of actuated spine cables. Figure 4.6 shows the results of these 90,000 trials. As can be seen, the best trial for the Achilles configuration, at 458.3 cm/min, is roughly three times as fast as NoFeet, SpiralsOnly, and ReducedSpirals after Monte Carlo, and is even faster than the evolved results for these other three cable configurations, at 807.4 cm/min. Figure 4.6 shows that the majority of the trials are concentrated below a threshold of 200 cm/min, which is much higher than the 50 cm/min threshold of SpiralsOnly and ReducedSpirals, and the 20 cm/min threshold of NoFeet. Also, the results in the upper half of this plot are not quite as sparse as they are for the other experiments. This indicates that, again, running more Monte Carlo trials would improve our evolutionary results. It is quite possible that the strong ground reaction forces generated with the help of these long, antagonistic pairs, contributes to this increase in speed.

Figure 4.7a show the progress of evolution for NoFeet, SpiralsOnly, ReducedSpirals, and Achilles. The first half of improvement occurs within the first 20 genera-


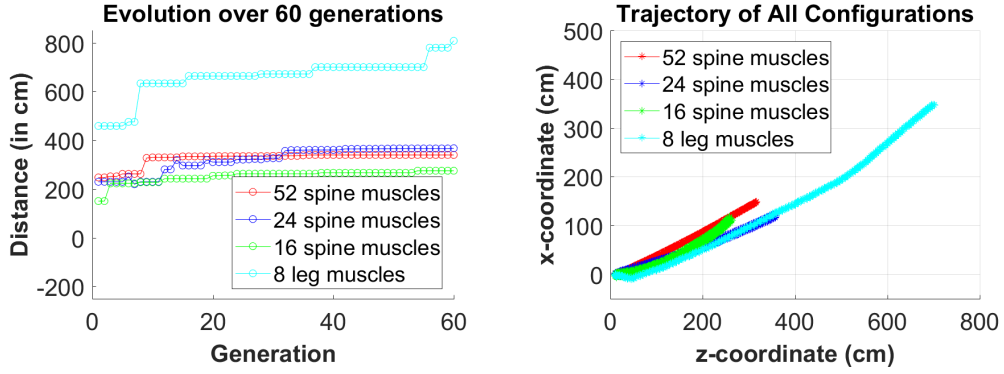


Figure 4.6: The initial 90,000 randomized Monte Carlo trials, each one minute in duration, for Achilles. For the results of the best 40 of these trials, see table 4.2 and figure 4.7a.

tions of evolution, in most cases leaving relatively minor improvements to the last 40 generations. This shows that our genetic algorithm converges quickly to the most advantageous solutions for a given morphology. ReducedSpirals both begins and ends the evolutionary stage as the slowest of all four actuator configurations, although during generations three through eleven it is tied with SpiralsOnly. Although SpiralsOnly and NoFeet start out close in performance, NoFeet overtakes SpiralsOnly as early as the tenth generation. Nevertheless, SpiralsOnly barely wins out over NoFeet by generation 30. Achilles starts out performing better than SpiralsOnly and ReducedSpirals, and therefore continues to outperform both throughout the entire evolution stage.

Figure 4.7b compares the resulting trajectories of the various actuator configurations. The paths traveled are the results after evolution, and the distances are listed



(a) A comparison of the progression of genetic evolution for NoFeet (red), SpiralsOnly (blue), ReducedSpirals (green), and Achilles (cyan). (b) The trajectories of NoFeet (red), SpiralsOnly (blue), ReducedSpirals (green), and Achilles (cyan), based on the center of mass of each robot.

Figure 4.7: The evolution process and the resulting trajectory of MountainGoat.

Table 4.2: Longest Distance Per Minute for Each Model

Model	Monte Carlo	Evolution
NoFeet	247.1 cm/min	339.4 cm/min
SpiralsOnly	229.4 cm/min	366.8 cm/min
ReducedSpirals	150.2 cm/min	274.7 cm/min
Achilles	458.3 cm/min	807.4 cm/min

in table 4.2. When plotted, some of the trajectories were rotated with respect to the origin in order to better compare distance. Note that all trajectories are, for the most part, straight. The distances, however, are very different. These different distances further support the concept, which we discussed in our previous work [14], that morphological design and control are highly coupled. A relatively minor reduction in the number of active cables from 24 to 16, for instance, caused a notable decrease in the speed of locomotion. A change in mode of actuation, from spine-driven to leg-driven locomotion, also resulted in an even more notable increase in the speed of locomotion.

Now to return to the question of whether reducing the amount of active cables in the spine would slightly or severely degrade performance. Reducing the number of spine actuators from 52 to 24 resulted in graceful degradation after Monte Carlo, but resulted in slightly improved performance after evolution. This improved performance after reducing the number of actuators was unexpected, and counter to the results found in [17]. Reducing the number of spine actuators to 16, however, showed a slightly greater degradation of performance after both Monte Carlo and evolutionary stages of learning. This is in line with results cited in [17]. Despite this slower performance, however, the robot was still able to travel a distance of almost three times its own body length in one minute, which is still respectable. Although using only 8 leg actuators resulted in substantially increased performance, this could be due to the employment a different mechanism of propulsion (i.e. legs instead spine).

4.2.2 Solution Space

Mirletz et. al., in their earlier analysis of solution space dimensionality of CPG networks, arrived at the following formula for the number of possible coupling weight and phase offset pairs that a morphology with repeating segments and bi-directional CPG couplings would be required to learn for productive motion [21]:

$$CPG \text{ couplings} = \frac{m(3m + 1)}{2} \quad (3)$$

With m representing the number of active cables, and hence number of active

Table 4.3: Number of parameter pairs learned for each model

Model	Active Cables	Max per Segment	Weight-Coupling Pairs
NoFeet	52	16	392
SpiralsOnly	24	8	100
ReducedSpirals	16	8	100
Achilles	8	2	7

CPG nodes, per segment. The configuration with 52 spine actuators, as shown in figure 1.2a, has 16 nodes per segment, while the configurations with 24 and 16 spine actuators each have 8 nodes per segment. the configuration with 8 leg actuators, Achilles, has only 2 nodes per leg. As seen in table 4.3, this equates to a larger solution space for NoFeet, which requires learning 392 coupling pairs. The other spine configurations, SpiralsOnly and ReducedSpirals, have smaller solution spaces, with 100 coupling pairs that need to be learned for productive locomotion. In contrast, Achilles has only 7 coupling pairs to be learned. Despite the fact that Achilles has extra parameters to be learned for the PD and impedance controllers, these only account for 5 more parameters, and thus this configuration has the smallest solution space of all four configurations.

4.3 Integration of Spine and Leg Control

Initial examination of the results shown in figure 4.7b might lead to the intuitive conclusion that pure leg actuation is preferable to pure spine actuation. Indeed the overwhelming majority of quadrupedal and bipedal robots take this approach. But

natural selection has found it advantageous for humans and other legged creatures to have both legs and a spine, and both of these structures have important roles in locomotion, as Gracovetsky, et. al. [12] and Folkertsma, et. al. [10], among many others, have theorized. It is only natural, then, to combine both legs and spine to see how speed of locomotion in MountainGoat is improved.




Figure 4.8: The initial 90,000 randomized Monte Carlo trials, each one minute in duration, for the model using both active spine and leg muscles.

As SpiralsOnly proved to be the fastest of the three spine-driven actuator configurations explored in 4.2.1, we chose to combine this configuration with that of Achilles for this experiment. We called this combined spine- and leg-actuated system AchillesSpirals. As with Achilles, we ran 90,000 Monte Carlo trials to allow for exploration over a solution space including PD and impedance controller parameters. Figure 4.8 shows the spread of the resulting 90,000 Monte Carlo trials. The longest distance

traveled during the Monte Carlo stage, as listed in table 4.4, was 787.0 cm/minute, and after genetic evolution the longest distance increased to 1189.2 cm/min.

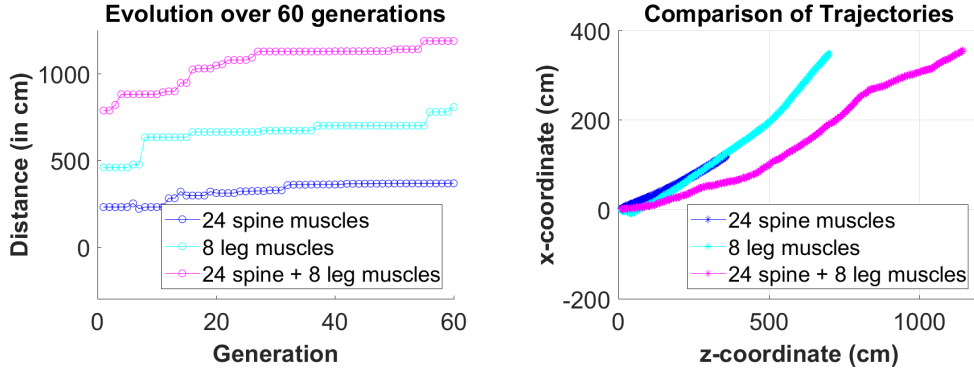
Even more so that for Achilles (shown in figure 4.6), the upper half of the Monte Carlo plot in figure 4.8 shows a denser concentration of trials, especially between the distances of 200 and 500. Again, this means that running even more than 90,000 Monte Carlo trials would potentially yield more trials with longer distances, which in turn would improve our evolutionary results.

Table 4.4: Longest Distance for Achilles and Number of Parameter Pairs

Model	Monte Carlo	Genetic Evolution	Active Cables	weight-coupling pairs
AchillesSpirals	787.0 cm/min	1189.2 cm/min	32	107

This evolutionary result shows an almost 50% improvement over Achilles, as listed in table 4.2. The total number of active cables in AchillesSpirals, also listed in table 4.4, is 32, and the total number of weight-coupling pairs that makes up the solution space is 107. This is still fewer actuated cables and, consequently, a smaller solution space than for NoFeet. But the resulting performance of AchillesSpirals is 3.5 times greater than NoFeet.

Figure 4.9a compares the process of evolution for AchillesSpirals with that of SpiralsOnly and Achilles. As can be seen, the evolution of AchillesSpirals starts and ends a faster speed than both SpiralsOnly and Achilles. As with the previous examples of genetic evolution shown above, most of the evolutionary work is done within the first 20



(a) Comparing the evolution of AchillesSpirals (magenta) over 60 generations to that of SpiralsOnly (blue) and Achilles (cyan). (b) Comparing the trajectory of AchillesSpirals (magenta) with those of SpiralsOnly (blue) and Achilles (red).

Figure 4.9: Evolution and Trajectory of SpiralsOnly, Achilles, and AchillesSpirals, to show improvement when both are combined.

generations, with only minor improvements made during the remaining 40 generations, showing convergence toward the best solution.

In figure 4.9b, the trajectory of AchillesSpirals's center of mass is compared with those of SpiralsOnly and Achilles. As in figure 4.7b, the trajectories of each actuator configuration are rotated about the origin in the plot for better comparison of distance and straightness of path. The trajectory of AchillesSpirals shows that the distance it traveled is roughly the amounts of SpiralsOnly and Achilles added together. Note, however, that the path of AchillesSpirals is not quite as smooth as those of SpiralsOnly and Achilles. One possible reason for this is that while there is nearest neighbor coupling for CPGs in the spine, as well as nearest neighbor coupling within each leg, this nearest neighbor coupling is not preserved between the legs and the spine. A way to reintegrate nearest neighbor coupling between the legs and the spine, which could be explored in future work, would be to create a hierarchy of CPGs along with a

more sophisticated feedback system for the higher level CPGs, as has been explored by Markin, et. al. [20]. A system such as this could lead to better coordination between the legs and the spine of the robot.

Chapter 5

Conclusions and Future Work

Our research explored the coupled aspects of morphological design of a spine-driven tensegrity quadruped, MountainGoat, and evaluation of the resulting CPG controlled locomotion in simulation. Each improvement to structural design increased the distance traveled by the robot. These results show that we have gained an understanding of the process of whole-body control, where the spine is central to locomotion and how extra support of the shoulders from the spine are necessary in order to lift the legs.

Future work could involve fine tuning aspects of the model such as the optimal number of vertebrae, distance between legs, shape of legs and vertebrae, stiffness and pretension of muscles, and arrangement of muscles. Using co-evolution of the robot's morphology and control would be an efficient means of exploring these aspects. Adding more torsion and the ability to bend the spine the the sagittal plane could be added, to help facilitate the lifting of legs. Knees could be added to the legs to aid locomotion over rough terrain, and these joints could be designed to become mechanically stiff in

the stance phase but become passively compliant in the swing phase.

Several active cable configurations for MountainGoat were explored, with each configuration reducing the solution space and improving speed over the original. The active spine configurations using 24 and 16 cables produced more productive locomotion than the configuration of 52 active cables, with the 24 cable configuration traveling the fastest of all configurations after evolution. The active leg configuration, with a total of 8 active cables, also showed improvement over the 52 cable spine configuration. Future work could include verification of these results with construction of actuated prototypes, to test the practicality of the various designs. Different gaits could also be explored in simulation, using gait-dependent offset pairs for individual legs.

Spine and Leg actuation were combined in an effort at whole-body control of MountainGoat, leading to a 50% speed increase over legs alone. Future work to improve the coordination between legs and spine would include a CPG hierarchy with a more robust feedback system for the higher level CPGs. Such a hierarchy would help smooth the trajectory of the robot's path.

It is interesting to note that morphological design and actuation strategy makes a notable difference, as over the course of morphological exploration presented in this thesis there was greater than 10X improvement in speed, even when reducing the total number of actuators. This shows that, when approached intelligently, efficiency of design and actuation can be achieved at the same time that performance is improved.

Bibliography

- [1] E Ackerman. Spot is boston dynamics nimble new quadruped robot. IEEE Spectrum, 2015.
- [2] Adrian Agogino, Vytas SunSpiral, and David Atkinson. Super Ball Bot - structures for planetary landing and exploration. NASA Innovative Advanced Concepts (NIAC) Program, Final Report, 2013.
- [3] Arvind Ananthanarayanan, Mojtaba Azadi, and Sangbae Kim. Towards a bio-inspired leg design for high-speed running. Bioinspiration & biomimetics, 7(4):046005, 2012.
- [4] R Adam Bilodeau, Edward L White, and Rebecca K Kramer. Monolithic fabrication of sensors and actuators in a soft robotic gripper. In Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, pages 2324–2329. IEEE, 2015.
- [5] Ken Caluwaerts, Jérémie Despraz, Atıl Işçen, Andrew P Sabelhaus, Jonathan Bruce, Benjamin Schrauwen, and Vytas SunSpiral. Design and control of com-

- pliant tensegrity robots through simulation and hardware validation. *Journal of The Royal Society Interface*, 11(98):20140520, 2014.
- [6] Jonas Degraeve, Ken Caluwaerts, Joni Dambre, and Francis Wyffels. Developing an embodied gait on a compliant quadrupedal robot. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 4486–4491. IEEE, 2015.
 - [7] JM Duperret, GD Kenneally, JL Pusey, and DE Koditschek. Towards a comparative measure of legged agility. In *Experimental Robotics*, pages 3–16. Springer, 2016.
 - [8] Daniel P Ferris, Gregory S Sawicki, and Monica A Daley. A physiologist’s perspective on robotic exoskeletons for human locomotion. *International Journal of Humanoid Robotics*, 4(03):507–528, 2007.
 - [9] Tom Flemons. The bones of tensegrity. http://www.intensiondesigns.com/bones_of_tensegrity, 2012.
 - [10] Gerrit A Folkertsma, Sangbae Kim, and Stefano Stramigioli. Parallel stiffness in a bounding quadruped with flexible spine. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2210–2215. IEEE, 2012.
 - [11] Sébastien Gay, José Santos-Victor, and Auke Ijspeert. Learning robot gait stability using neural networks as sensory feedback function for central pattern generators.

- In Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on, pages 194–201. IEEE, 2013.
- [12] S Gracovetsky. An hypothesis for the role of the spine in human locomotion: a challenge to current thinking. *Journal of biomedical engineering*, 7(3):205–216, 1985.
 - [13] Sten Grillner, Alexander Kozlov, Paolo Dario, Cesare Stefanini, Arianna Menciassi, Anders Lansner, and Jeanette Hellgren Kotaleski. Modeling a vertebrate motor system: pattern generation, steering and control of body orientation. *Progress in brain research*, 165:221–234, 2007.
 - [14] D Hustig-Schultz, V SunSpiral, and M Teodorescu. Morphological Design for Controlled Tensegrity Quadruped Locomotion. In *Proceedings of 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4714–4719. IEEE, Oct. 2016.
 - [15] Atil Iscen, Adrian Agogino, Vytas SunSpiral, and Kagan Tumer. Robust distributed control of rolling tensegrity robot. In *The Autonomous Robots and Multirobot Systems (ARMS) workshop at AAMAS 2013*, 2013.
 - [16] Atil Iscen, Ken Caluwaerts, Jonathan Bruce, Adrian Agogino, Vytas SunSpiral, and Kagan Tumer. Learning tensegrity locomotion using open-loop control signals and coevolutionary algorithms. *Artificial life*, 2015.
 - [17] S Lessard, J Bruce, A Agogino, V SunSpiral, and M Teodorescu. Robust Monte

- Carlo Control Policies to Maneuver Tensegrity Robots out of Obstacles. In Proceedings of Autonomous Robots and Multirobot Systems (ARMS), May 2015, Istanbul, Turkey, 2015.
- [18] S Lessard, D Castro, W Asper, S Chopra, L Baltaxe-Admony, M Teodorescu, V SunSpiral, and A Agogino. A bio-inspired tensegrity manipulator with multi-dof, structurally compliant joints. In Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2016.
- [19] Stephen Levin. The tensegrity-truss as a model for spine mechanics: Biotensegrity. *Journal of Mechanics in Medicine and Biology*, 2:375–388, 2002.
- [20] Sergey N Markin, Alexander N Klishko, Natalia A Shevtsova, Michel A Lemay, Boris I Prilutsky, and Ilya A Rybak. A neuromechanical model of spinal control of locomotion. In *Neuromechanical Modeling of Posture and Locomotion*, pages 21–65. Springer, 2016.
- [21] B Mirletz, In-Won Park, Thomas E Flemons, Adrian K Agogino, Roger D Quinn, and Vytas SunSpiral. Design and control of modular spine-like tensegrity structures. In *The 6th World Conference of the International Association for Structural Control and Monitoring (6WCSCM)*, 2014.
- [22] Brian T Mirletz, Perry Bhandal, Ryan D Adams, Adrian K Agogino, Roger D Quinn, and Vytas SunSpiral. Goal-Directed CPG-Based Control for Tensegrity

- Spines with Many Degrees of Freedom Traversing Irregular Terrain. *Soft Robotics*, 2(4):165–176, 2015.
- [23] Brian T Mirletz, In-Won Park, Roger D Quinn, and Vytas SunSpiral. Towards bridging the reality gap between tensegrity simulation and robotic hardware. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 5357–5363. IEEE, 2015.
- [24] Brian T Mirletz, Roger D Quinn, and Vytas SunSpiral. CPGs for Adaptive Control of Spine-like Tensegrity Structures. In *Proceedings of 2015 International Conference on Robotics and Automation (ICRA2015) Workshop on Central Pattern Generators for Locomotion Control: Pros, Cons & Alternatives*, 2015.
- [25] Brian Tietz Mirletz. Adaptive Central Pattern Generators for Control OF Tensegrity Spines with Many Degrees of Freedom. PhD thesis, Case Western Reserve University, 2016.
- [26] Mario Mulansky and Karsten Ahnert. Odeint library. *Scholarpedia*, 9(12):32342, 2014.
- [27] Ryuma Niiyama, Carine Rognon, and Yasuo Kuniyoshi. Printable pneumatic artificial muscles for anatomy-based humanoid robots. In *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*, pages 401–406. IEEE, 2015.

- [28] O Orki, A Ayali, O Shai, and U Ben-Hanan. Modeling of caterpillar crawl using novel tensegrity structures. *Bioinspiration & biomimetics*, 7(4):046006, 2012.
- [29] Chandana Paul, John W Roberts, Hod Lipson, and Francisco J Valero Cuevas. Gait production in a tensegrity based robot. In *Advanced Robotics, 2005. ICAR'05. Proceedings., 12th International Conference on*, pages 216–222. IEEE, 2005.
- [30] Ron Pelrine, Roy D Kornbluh, Qibing Pei, Scott Stanford, Seajin Oh, Joseph Eckerle, Robert J Full, Marcus A Rosenthal, and Kenneth Meijer. Dielectric elastomer artificial muscle actuators: toward biomimetic motion. In *SPIE's 9th Annual International Symposium on Smart Structures and Materials*, pages 126–137. International Society for Optics and Photonics, 2002.
- [31] Michael T Petralia and Robert J Wood. Fabrication and analysis of dielectric-elastomer minimum-energy structures for highly-deformable soft robotic systems. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 2357–2363. IEEE, 2010.
- [32] Panagiotis Polygerinos, Stacey Lyne, Zheng Wang, Luis Fernando Nicolini, Bobak Mosadegh, George M Whitesides, and Conor J Walsh. Towards a soft pneumatic glove for hand rehabilitation. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1512–1517. IEEE, 2013.
- [33] Marc Raibert, Kevin Blankespoor, Gabriel Nelson, Rob Playter, and TB Team.

- Bigdog, the rough-terrain quadruped robot. In Proceedings of the 17th World Congress, volume 17, pages 10822–10825, 2008.
- [34] Ludovic Righetti, Jonas Buchli, and Auke Jan Ijspeert. Dynamic hebbian learning in adaptive frequency oscillators. *Physica D: Nonlinear Phenomena*, 216(2):269–281, 2006.
- [35] Andrew P Sabelhaus, Hao Ji, Patrick Hylton, Yakshu Madaan, ChanWoo Yang, Alice M Agogino, Jeffrey Friesen, and Vytas SunSpiral. Mechanism design and simulation of the ultra spine: A tensegrity robot. In ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, pages V05AT08A059–V05AT08A059. American Society of Mechanical Engineers, 2015.
- [36] Sangok Seok, Albert Wang, Meng Yee Michael Chuah, Dong Jin Hyun, Jongwoo Lee, David M Otten, Jeffrey H Lang, and Sangbae Kim. Design principles for energy-efficient legged locomotion and implementation on the mit cheetah robot. *Mechatronics, IEEE/ASME Transactions on*, 20(3):1117–1129, 2015.
- [37] Gurucharan Singh, Vijay P Agrawal, M Basavarajappa, and M Manohar. Gastrocnemius-achilles tendon-axis: A human anatomical variation. *International Journal of Biomedical and Advance Research*, 3(8):653–655, 2012.
- [38] Robert E Skelton, Rajesh Adhikari, J-P Pinaud, Waileung Chan, and JW Helton. An introduction to the mechanics of tensegrity structures. In *Decision and Control*,

2001. Proceedings of the 40th IEEE Conference on, volume 5, pages 4254–4259. IEEE, 2001.
- [39] Kenneth Snelson. Continuous tension, discontinuous compression structures. united states patent 3169611, February 1965.
- [40] Brian R Tietz, Ross W Carnahan, Richard J Bachmann, Roger D Quinn, and Vytas SunSpiral. Tetraspine: Robust terrain handling on a tensegrity robot using central pattern generators. In AIM, pages 261–267, 2013.
- [41] T Umedachi, V Vikas, and BA Trimmer. Softworms: the design and control of non-pneumatic, 3d-printed, deformable robots. *Bioinspiration & biomimetics*, 11(2):025001, 2016.
- [42] Alexander Lawrence Xydes. Simulating DuCTT and optimizing control for DuCTT with machine learning. PhD thesis, University of California, San Diego, 2015.
- [43] Qian Zhao, Hidenobu Sumioka, Kohei Nakajima, Xiaoxiang Yu, and Rolf Pfeifer. Spine as an engine: effect of spine morphology on spine-driven quadruped locomotion. *Advanced Robotics*, 28(6):367–378, 2014.