

©Copyright 2025

Joseph G. Sullivan

Toward Energy-Efficient Actuation of Legged Locomotion Using
Handed Shearing Auxetic Parallel Elastic Structures

Joseph G. Sullivan

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Sam Burden, Chair

Blake Hannaford

Kim Ingraham

Jeffrey Lipton

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

Abstract

Toward Energy-Efficient Actuation of Legged Locomotion Using Handed Shearing Auxetic Parallel Elastic Structures

Joseph G. Sullivan

Chair of the Supervisory Committee:
Associate Professor Sam Burden
Electrical and Computer Engineering

This dissertation explores the potential of Handed Shearing Auxetic (HSA) structures as components in energy-efficient robotic actuators. We present a novel parallel elastic actuator that integrates an HSA with a quasi-direct drive motor, combining passive compliance and static braking in a compact and mechanically efficient design. To characterize the actuator’s nonlinear viscoelastic behavior, we develop a structured modeling framework grounded in Lagrangian mechanics, using convex elastic and dissipation potentials. A variational loss based on the Euler–Lagrange residual enables tractable system identification from trajectory data using only motor telemetry.

We validate this approach on a vertically hopping monopod robot, where the learned models are integrated into a trajectory optimization framework. Experiments demonstrate improved electrical efficiency during hopping and load-bearing tasks, particularly under optimized control. These results suggest that HSAs can contribute meaningfully to compliant actuation and motivate further research into their use in more general legged locomotion systems.

TABLE OF CONTENTS

	Page
List of Figures	iii
Glossary	v
Chapter 1: Introduction	1
Chapter 2: Background and Related Work	3
2.1 Compliant Actuation in Legged Robots	3
2.2 Properties and Applications of Handed Shearing Auxetics	5
2.3 Kinematic and Dynamic Models of HSA Robots	6
2.4 Related Works of Physics-Informed Machine Learning	7
Chapter 3: Design of and Evaluation the HSA-PEA Hopping Robot	9
3.1 The Spring-Mass Oscillator Template	10
3.2 HSA Selection and Fabrication	11
3.3 HSA Parallel Compliance Mechanism	13
3.4 Hardware and Electronics	15
3.5 Static Braking Tests	17
3.6 Model-Free Hopping Experiments	17
3.7 Results	19
3.8 Conclusion	22
Chapter 4: Data-Driven Viscoelastic Modeling	24
4.1 Lagrange-D'Alembert Generalized Action Principle	24
4.2 Elastic Potentials	26
4.3 Dissipation Potentials	27
4.4 Viscoelastic HSA Model	28
4.5 Data-Driven Model Identification	29
4.6 Viscoelastic Model Validation	32
4.7 Results	33

4.8	Discussion	37
4.9	Limitations	37
4.10	Variational Interpretation and Future Extensions	39
Chapter 5:	Optimized Hopping Experiments	41
5.1	Problem Formulation	41
5.2	Data Collection and Statistical Analysis	42
5.3	Results	44
5.4	Discussion	47
Chapter 6:	Conclusion	48
6.1	Summary of Contributions	48
6.2	Limitations	49
6.3	Future Directions	50
6.4	Closing Remarks	51
Bibliography	52

LIST OF FIGURES

Figure Number	Page
3.1 System Overview. We use Handed Shearing Auxetics (HSA) as a combined Spring and Break mechanism in a hopping leg to improve system efficiency in hopping and standing.	10
3.2 HSA average stiffness over five-centimeters stroke is shown for different twist angles. Data was collected on an instron UTM. For large twist angles (red region), the HSA jams against the inner cylinder, dramatically increasing its stiffness.	12
3.3 Exploded view of the HSA assembly, illustrating the twisting mechanism at its base. The HSA is fixed at its right end to a gear carrier, which is driven by a compact servomotor (“Twist Motor”). The left end is coupled to the foot via a 30 cm steel rod passing through a linear bearing embedded in the base. A shaft collar and thrust bearings secure the gear assembly to the cart (not shown). Inside the HSA, a rigid cylindrical insert (occluded in orange) enables the braking mechanism by limiting auxetic deformation.	13
3.4 The HSA-PEA Hopping robot is shown in two poses demonstrating auxetic deformation of the parallel HSA.	15
3.5 Electrical system block diagram. Solid arrows represent power interconnections, dashed arrows denote data and control lines. The actuator is controlled by an off-board microcomputer integrated with a full-duplex controller area network (CAN) interface.	16
3.6 Comparison of the electrical power drawn vs. the force blocked by the leg motor compared to the HSA twist motor. The power consumed by the primary motor (squares) is fit by a quadratic function of the blocked force (purple dotted line). The power consumed by the HSA twist motor (diamonds) follows a linear trend (gold dashed line). The solid gold line predicts the power that would be consumed by the twist servo at 1:8 gear reduction.	18
3.7 Part (a): Hopping cost-of-transport with and without an HSA compared to results from the SPEAR hopping robot. Part (b): Changes to the mechanical and thermal components of COT (with an HSA minus without an HSA), shown in <i>absolute value</i> , i.e. the blue curve is the <i>decrease</i> in the thermal cost, the red curve is the <i>increase</i> in mechanical cost. The reduced hopping COT in (a) is due to reduced thermal cost, while the mechanical cost increases, suggesting energy is dissipated by the HSA.	20

4.1	Illustration of the integral quadrature method for evaluating model error. Top: Measured HSA length over a single hop (purple) with quadrature points (green). Bottom: Residuals of the Euler-Lagrange equation (green) and weighted residuals (purple). A systematic error appears at landing due to missing ground impulse data but has reduced influence due to low quadrature weight.	33
4.2	Model-predicted elastic (left) and dissipative (right) forces across all test hops. Each trace represents a different HSA twist angle θ . Nonlinear behavior is evident in both force types: elastic force shows varying curvature, and damping force exhibits configuration and rate dependent hysteresis.	34
4.3	Model-predicted viscoelastic force across all test hops. Each trace shows the sum of the contributions from the elastic and dissipation potential at a different HSA twist angle θ	34
4.4	Heatmap of the effective spring rate for the elastic potential in model (1). The dashed line shows the zero-force length at each twist angle. The spring rate is strictly positive in accordance with strong convexity.	35
5.1	Scatter plot of Cost of Transport versus hopping height for all trials. Each condition is color-coded and partially transparent to visualize density. Overlap across groups and high variance motivate the need for regression to isolate effects of optimization and the HSA.	44
5.2	Boxplots comparing Cost of Transport and hopping height across the four experimental conditions: HSA vs. no HSA and optimized vs. not optimized. While both optimization and the HSA reduce mean COT, large and unequal variances complicate direct comparisons. Differences in hopping height across groups motivate controlling for this covariate in regression analysis.	45
5.3	Left: Bootstrapped sample means ($n = 40, 10,000$ resamples) of COT versus hopping height for each experimental group. Regression lines highlight trends. Right: Predicted COT at the median hopping height (5.43 cm) with 95% confidence intervals. Both optimization and the HSA significantly reduce COT, and the benefit of the HSA is amplified under optimized trajectories.	46

GLOSSARY

HSA (HANDED SHEARING AUXETIC): : A periodic structure with auxetic geometry that expands diametrically under torsion and transverse strain, featured as a compliant element in soft robots.

COT (COST OF TRANSPORT): A dimensionless measure of energy efficiency during locomotion, defined as energy per unit weight per unit distance.

PEA (PARALLEL ELASTIC ACTUATOR): An actuator configuration where a compliant element is placed in parallel with a motor, reducing motor load.

VISCOELASTICITY: : A material property that characterized by elastic (reversible) deformation and rate dependent hysteresis (viscous dissipation).

ACKNOWLEDGMENTS

The author wishes to express his gratitude to Dr. Sam Burden and Dr. Ian Good, two scholars who are exceptional in their capacity for kindness and empathy.

DEDICATION

to my dear parents, Bryan and Karen

Chapter 1

INTRODUCTION

Legged robots have demonstrated exceptional versatility in environments where wheeled and tracked systems struggle, making them well-suited for inspection, exploration, and service applications. However, their range and endurance remain limited, in large part due to poor energy efficiency [3]. Compared to animals and wheeled machines, legged systems consistently exhibit higher cost of transport (COT)—a fundamental metric quantifying the energy required to move a unit mass over a unit distance [17]. In addition to dynamic inefficiencies during locomotion, legged robots frequently expend energy simply to maintain posture, posing a challenge for long-duration tasks such as stationary sensing or manipulation.

One of the most successful strategies to reduce COT has been the integration of mechanical compliance into legged robot actuators [13]. Springs can store and return mechanical energy during periodic motion, reducing the electrical burden on motors. Series elastic actuators and parallel elastic actuators have been employed to exploit this effect, with examples ranging from bipedal machines like Cassie [29] to quadrupeds like ANYmal [14]. Additional mechanisms such as clutches and locks have been used to decouple compliance during specific phases of motion [20] or to reduce static power consumption [26]. While these designs have proven effective, they often rely on discrete mechanisms and conventional linear springs.

Recent advances in materials science and additive manufacturing have opened the door to new actuator designs based on architected metamaterials [23, 43]—engineered structures with tailored geometry and nonlinear mechanics. These systems blur the line between rigid and soft robotics [31], combining rigid-body motion with deformation-based compliance and dissipation. One such class of materials, Handed Shearing Auxetics (HSAs) [18], exhibits unusual mechanical behavior under coupled torsional and axial loads, enabling functions

such as jamming-based braking and programmable stiffness. Embedding such structures within robotic actuators presents an opportunity to unify compliance and locking functionality, potentially improving both dynamic and static efficiency.

However, these benefits come with new challenges. Architected structures often exhibit nonlinear, configuration-dependent viscoelastic [6, 42] behavior that lies outside the scope of rigid-body mechanics and simple models of elastic energy. Capturing their behavior requires moving beyond traditional design heuristics and toward systematic, data-driven modeling approaches [4]. While deep learning methods have been used to model soft robotic dynamics, their black-box nature introduces issues of interpretability and data efficiency. As an alternative, we explore a physics-informed framework grounded in analytical mechanics, enabling interpretable and statistically testable models of viscoelastic structures embedded in actuators.

Contributions. This dissertation presents a unified experimental and modeling framework for a novel parallel elastic actuator based on Handed Shearing Auxetics. The primary contributions are:

- We present the design, fabrication, and experimental evaluation of an HSA-based actuator integrated into a monopod hopping robot, demonstrating improved energy efficiency in both dynamic hopping and static posture maintenance.
- We develop a variational modeling framework that combines rigid-body dynamics with convex potential representations of elasticity and dissipation, fitted from data using a convex regression based on the Euler–Lagrange residual.
- We integrate the learned models into a trajectory optimization pipeline, and show that model-informed control strategies further improve energy efficiency over heuristic baselines.
- We perform statistical validation of model structure and generalization ability, confirming that nonlinear elasticity and dissipation play significant roles in the actuator’s behavior.

Chapter 2

BACKGROUND AND RELATED WORK

Legged robots, particularly those incorporating compliant actuation, have motivated a diverse body of research across mechanical design, materials science, and data-driven modeling. This chapter reviews key developments in each of these domains to contextualize our work. We begin by comparing series and parallel elastic actuation strategies, highlighting their tradeoffs in dynamic and static tasks. We then introduce Handed Shearing Auxetics (HSAs), a class of architected metamaterials whose geometric mechanics support both spring-like and brake-like behavior. To support integration of HSAs into dynamic systems, we survey relevant models from soft robotics and viscoelastic materials, identifying gaps in existing approaches. Finally, we review recent advances in physics-informed machine learning that enable interpretable modeling of nonlinear mechanical systems from data.

2.1 Compliant Actuation in Legged Robots

Despite their mobility advantages, electrically actuated legged robots suffer from poor energy efficiency relative to wheeled or tracked systems. Several key factors contribute to this limitation. First, legged locomotion often demands high torques at low speeds—conditions under which electric motors generate significant thermal losses due to Joule heating [17]. Second, legged systems struggle to recycle energy effectively: while biological locomotion exploits elastic energy storage and restitution, robotic actuators tend to dissipate impact energy rather than recover it [3]. Third, during static postures or low-speed tasks, electric actuators must supply continuous torque to support the robot’s weight [26], resulting in additional power draw and overheating risk.

To address these challenges, researchers have explored various actuation architectures that incorporate mechanical compliance and passive elements. Two major paradigms have emerged: series elastic actuators (SEAs) and parallel elastic actuators (PEAs). On this

topic, Yesilevsky et al [41] provided a comprehensive theoretical analysis of SEAs vs PEAs in hopping tasks.

In the general sense of legged-locomotion SEAs are the more common approach. Recent examples include ANYmal [14], Salto [10], Atrias [12], and the Michigan Prosthesis [1]. Within the SEA paradigm, a compliant spring is placed in series between the motor and the load, allowing energy storage during motion and shock absorption at impact. Because SEAs typically use high-reduction gearboxes to improve torque density, they are particularly effective for tasks involving low-speed cyclic motions. However, they introduce complexity in control due to phase lag and additional dynamics from the gearbox and spring [13]. SEAs have been shown to improve hopping efficiency when regeneration is ignored, but they are less effective in regimes where energy must be dissipated or where high-fidelity torque control is required [29].

PEAs, by contrast, place the spring in parallel with the motor, often using a low-reduction, quasi-direct-drive actuator. This configuration allows the motor and spring to share the load, offloading energy demands from the motor during cyclical motion. Legged robots featuring PEAs include a progenitor of Cheetah [7] and a recent descendant [19], as well as BioBiped [34], and the SPEAR hopper [40]. PEAs offer improved mechanical efficiency and simplicity over SEAs, particularly when the motor must track high-speed trajectories or handle variable loads [33]. However, PEAs suffer from poor thermal efficiency in static conditions and increased impact losses due to higher reflected inertia [17]. This has prompted investigation into design trade-offs and hybrid actuator configurations that combine the advantages of both paradigms.

To compensate for these limitations, various locking mechanisms have been explored. Cams, ratchets, and clutches can passively redirect the energy of swinging legs by locking a joint entirely, thereby eliminating some of the negative motor work. These strategies can be used in combination with compliance to selectively engage springs with the actuator, improving performance during both dynamic and static tasks [27]. While effective, each added component increases complexity, weight, and design overhead. This motivates interest in multi-functional mechanisms that can unify spring-like and brake-like behavior within a single structure.

2.2 *Properties and Applications of Handed Shearing Auxetics*

Handed Shearing Auxetics (HSAs) are a class of architected metamaterials that exhibit unusual mechanical responses under torsional loading [18]. Unlike traditional materials, which typically narrow under axial extension, HSAs expand in diameter as they elongate—a behavior linked to their chiral, helical geometry. This auxetic response allows HSAs to couple axial stretch with radial and shear deformation, enabling programmable stiffness and nonlinear energy storage [8]. Because their behavior is geometry-driven rather than material-dependent, HSAs can be easily fabricated using modern additive manufacturing tools, and scaled or tuned for different applications. They have been explored in a wide range of robotic systems, from soft grippers [5] and robot arms [39], and even a quasi-static quadrupedal robot [16].

In the context of legged robotics, HSAs offer a compelling opportunity to combine parallel elasticity and passive braking into a single actuator component. When embedded in a prismatic leg mechanism, an HSA can behave like a spring under axial compression or extension. If the auxetic deformation is constrained—for example, by twisting the structure against an inner insert—the HSA can “lock” into a jammed configuration, providing high resistance to further motion. This behavior is analogous to a capstan brake and enables low-power static load support, a feature particularly useful in robots that must pause or brace during extended tasks.

However, prior applications of HSAs in robotics has focused on quasi-static behavior, particularly in the small-strain regime where the elastic properties of HSAs has been successfully modeled [37, 38]. In contrast, the behavior of HSAs under high-speed and large-strain dynamic loading—such as those encountered during jumping or hopping—remains unexplored.

This work aims to fill that gap by integrating a 3D-printed HSA into a parallel elastic leg actuator, evaluating its behavior under cyclic loads, and modeling its nonlinear viscoelastic response. At the same time, the leverage the nonlinear auxetic trajectory to lock the leg joint, creating a passive braking effect. This locking mechanism enables the robot to support static loads with minimal power. As we will show, this results in a single component that

combines the functionality of a spring and a brake—a “spring-brake”—with applications in both dynamic and static tasks.

2.3 Kinematic and Dynamic Models of HSA Robots

While architected metamaterials such as HSAs offer promising mechanical properties for robotic actuation, their integration into dynamic systems presents a number of modeling challenges. In particular, HSAs do not conform to the assumptions of classical linear elasticity: their stiffness varies significantly with strain, and they exhibit strong hysteretic effects over long timescales. To effectively control and optimize systems incorporating HSAs, it is necessary to develop models that capture these nonlinear behaviors while remaining tractable and interpretable.

Sustained interest in soft robotics has led to the development of several reduced-order models to describe the kinematic and elastic properties of rod-like robots [4]. Among the most successful is the 3D Cosserat rod model, which approximates rod-like structures with a virtual backbone characterized by constant stiffness. This approach was later extended to capture piecewise-constant curvature (PCC), and further generalized by Renda et al. [30] into the Geometric Variable-Strain (GVS) model, which allows curvature to vary continuously along the backbone. Renda’s method, which represents elastic behavior as a linear combination of basis potentials, is conceptually related to the framework adopted in this work.

While the PCC and GVS approaches provide effective representations of kinematic and elastic behavior, the dissipation properties of HSAs have received far less attention. Quasi-static force characterization experiments show that HSAs often exhibit wide hysteresis loops across a range of materials [18]. Many HSAs in the robotics literature are printed from polyurethane resins such as FPU, TPU, and EPU using digital light synthesis (DLS). Several investigations into the mechanical properties of DLS-printed polyurethanes report significant viscous dissipation [25], including nonlinear, rate-dependent hysteresis effects in EPU [11]. The interactions between material dissipation and the geometric deformation of HSAs are likely to complicate this picture further. As a result, reduced-order, data-driven models may be necessary to accurately describe damping behavior in printed HSAs, particularly

when used in dynamic regimes.

Few studies to date have attempted to develop viscoelastic models of HSAs in robotic systems. Stölzle et al. [38] modeled a soft robotic arm actuated by four parallel HSAs using the PCC framework. They identified the elastic terms of a parametric model by minimizing the residuals of the Euler–Lagrange equations but did not estimate damping parameters—instead compensating for dissipation within the control policy. This approach proved effective in low-strain, quasi-static conditions where geometric nonlinearity was limited.

The actuator studied in this work departs significantly from that of Stölzle et al. Here, the HSA is mounted in parallel with a low-reduction actuator and mechanically constrained to eliminate bending strain. It operates in a high-strain, high-speed regime, mounted to a half-pantograph leg through rigid ABS components. In this domain, viscoelastic behavior is likely to be highly nonlinear and configuration-dependent. To model its behavior, we adopt a structured, data-driven approach that preserves the variational foundations of mechanics. Specifically, we formulate the actuator’s elastic and dissipative properties as convex potentials—a technique commonly used in rheological modeling and generalized finite element methods. The net force is expressed as the differential of these convex potential functions with respect to the relevant kinematic quantities.

This modeling approach balances expressive power with physical interpretability. By grounding the model in analytical mechanics, we retain direct access to energy and power quantities, enabling integration with trajectory optimization frameworks and supporting statistically principled system identification. In the chapters that follow, we demonstrate how this framework enables both predictive modeling and energy-efficient control of an HSA-driven hopping robot.

2.4 Related Works of Physics-Informed Machine Learning

In recent years, physics-informed machine learning has emerged as a powerful framework for discovering dynamical systems from data while preserving the underlying physical structure. These methods leverage priors from classical mechanics—such as conservation laws and variational principles—to improve data efficiency, interpretability, and generalization,

particularly in systems with complex or partially unknown dynamics.

Among the most influential approaches is the Sparse Identification of Nonlinear Dynamics (SINDy) framework [2, 36], which casts model identification as a sparse regression problem over a library of nonlinear functions. SINDy has proven effective in a wide range of systems, but it is typically formulated for first-order differential equations. As a result, it often requires finite-difference estimates of derivatives, which can be noisy or inaccurate for second-order mechanical systems—particularly under nonsmooth dynamics. Integral formulations of the residual have been proposed to mitigate these issues [?], an idea that closely relates to the variational approach adopted in this work.

Another line of research incorporates physical structure directly into neural network architectures. Hamiltonian Neural Networks (HNNs) [9] and Lagrangian Neural Networks (LNNs) [22] constrain model outputs to satisfy geometric priors from analytical mechanics. HNNs are trained to reproduce the system’s vector field as the gradient of a learned Hamiltonian, assuming conservative dynamics. LNNs take a similar approach using the Euler–Lagrange equations, learning both kinetic and potential energy functions from data. While these models are physically grounded, they generally rely on dense network architectures and do not account for non-conservative forces such as dissipation.

Recent extensions have sought to bridge this gap. Sharma et al. [35] introduced a reduced-order modeling framework that incorporates dissipation by identifying low-dimensional operators from data. Their method formulates operator fitting as a constrained least-squares problem, grounded in the Lagrangian framework. However, their elastic and damping forces are based on constant-curvature potentials and acceleration-coordinate formulations, which can obscure the individual contributions of elasticity and damping in complex systems.

In summary, physics-informed learning methods have shown great promise for modeling mechanical systems, but many existing approaches are either restricted to conservative dynamics or constrained by black-box architectures. The modeling framework presented in this work combines the interpretability of analytical mechanics with the flexibility of convex regression, enabling tractable system identification of viscoelastic structures with physically meaningful components.

Chapter 3

DESIGN OF AND EVALUATION THE HSA-PEA HOPPING ROBOT

This chapter presents the design of a parallel elastic leg actuator based on a 3D-printed Handed Shearing Auxetic (HSA) structure. The work described here was conducted in collaboration with Dr. Ian Good, Dr. Samuel A. Burden, and Dr. Jeffrey Lipton, and is currently under review with *IEEE Robotics and Automation Letters* (RAL). The actuator prototype, referred to as a "spring-brake" mechanism, leverages the nonlinear mechanical properties of the HSA to achieve both elastic energy storage during dynamic hopping and passive load support during static tasks.

The experiments and results in this chapter provide the physical foundation for the modeling and control framework developed in Chapter 4. In particular, the dynamic hopping tests described here expose complex viscoelastic behaviors—such as configuration-dependent stiffness and dissipation—that motivate the need for data-driven system identification. These behaviors are later quantified through a variational modeling approach, and the resulting models are used in Chapter 5 to generate optimized hopping trajectories. Together, these chapters form a cohesive narrative from hardware design, to model development, to performance-aware control of HSA-actuated legged robots.




Figure 3.1: System Overview. We use Handed Shearing Auxetics (HSA) as a combined Spring and Break mechanism in a hopping leg to improve system efficiency in hopping and standing.

3.1 The Spring-Mass Oscillator Template

In this chapter presents the design of a hopping robot capable of hopping at 5cm height at frequency $f > 2$ Hz. Our approach combines a low-reduction motor with an HSA in parallel. We selected a quasi direct-drive motor for its favorable ratio of torque to reflected inertia and ample regeneration capability.

To guide the design of the HSA, we approximated the stance-phase dynamics of the robot using a spring-mass oscillator model:

$$m\ddot{x} + K(x - x_0) + mg = 0,$$

where $m = 1.3$ Kg is the total system mass, x represents the leg length, K is the effective

stiffness, and g is the gravitational acceleration.

We estimated the required spring constant K by equating the mechanical energy at the lowest point of stance with the potential energy needed for a hop of height h :

$$mgh = \frac{1}{2}K\Delta x^2 - mg\Delta x.$$

Rearranging for K ,

$$K = 2mg(h + \Delta x)\Delta x^{-2}.$$

Setting $\Delta x = h = 5$ cm gives an effective spring rate of $K \approx 1 \frac{\text{kN}}{\text{m}}$.

To verify this stiffness yields a hopping frequency above 2 Hz, we calculate the stance and flight durations:

$$t_F = 2\sqrt{\frac{2h}{g}}, \quad t_S(m, K) \approx \pi\sqrt{\frac{m}{K}}$$

and the total hopping frequency:

$$f(m, K) \approx \left(2\sqrt{\frac{2h}{g}} + \pi\sqrt{\frac{m}{K}} \right)^{-1}. \quad (3.1)$$

Substituting $K = 1 \frac{\text{kN}}{\text{m}}$, we find:

$$f(1.3, 1000) \approx 3.2 \text{ Hz}.$$

This analysis yields the following actuator design requirements:

- Stroke length: $\Delta x > 5$ cm
- Effective Stiffness: $K \approx 1000 \frac{\text{N}}{\text{m}}$
- Load capacity: $F_{\max} \geq 1.5 \times K \times 5 \text{ cm} = 75 \text{ N}$

3.2 HSA Selection and Fabrication

We selected an HSA configured at 0% of its auxetic trajectory to maximize stroke length and minimize stiffness, following design guidelines from prior studies [8, 39]. An 8-row, 3-column HSA geometry was chosen to balance jamming capability with stroke range. The structure was printed in FPU50 using a Carbon M1 DLS printer per manufacturer specifications.

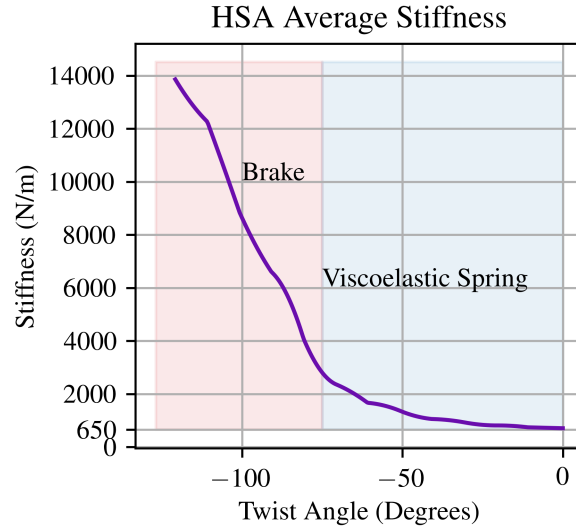


Figure 3.2: HSA average stiffness over five-centimeters stroke is shown for different twist angles. Data was collected on an instron UTM. For large twist angles (red region), the HSA jams against the inner cylinder, dramatically increasing its stiffness.

The HSA was made from FPU50 on the Carbon M1 DPL printer following manufactures best practices. FPU50 HSAs can easily support cyclic loads of 50% strain so the HSA was scaled to a length of 100 mm, with an additional 16 mm added for mounting, resulting in a 3D printed part that is 116mm long and 34mm in diameter. The minimum cross section in the design is 36 mm^2 , and with the assumed loading conditions, we see internal stresses of 2.1 MPa, which is well below UTS (25 MPa). Characterization of the printed HSAs were performed on an Instron UTM, sweeping twist angles from -150° to 120° in 10° increments, and displacement from 0 to 5cm in 5mm increments. Under positive twist, the average spring constant reached $912 \frac{\text{N}}{\text{m}}$, with peak stiffness exceeding $16 \frac{\text{kN}}{\text{m}}$ —a $21\times$ increase over the minimum. These results confirm that the fabricated HSA meets all design specifications.

3.3 HSA Parallel Compliance Mechanism




Figure 3.3: Exploded view of the HSA assembly, illustrating the twisting mechanism at its base. The HSA is fixed at its right end to a gear carrier, which is driven by a compact servomotor (“Twist Motor”). The left end is coupled to the foot via a 30 cm steel rod passing through a linear bearing embedded in the base. A shaft collar and thrust bearings secure the gear assembly to the cart (not shown). Inside the HSA, a rigid cylindrical insert (occluded in orange) enables the braking mechanism by limiting auxetic deformation.

The compliance mechanism integrates the HSA in parallel with a QDD100 motor (1:6 reduction), with leg extension x related to HSA length l by $l = x - 30$ cm. This arises from a closed-loop pantograph leg geometry involving four links (Table 3.3). Figure 3.3 shows the HSA assembly. The HSA is mounted at the base to a gear carrier driven by a small twist motor. The upper end is connected to the foot through a steel rod constrained by a linear bearing. A rigid cylindrical insert inside the HSA facilitates the braking behavior by restricting auxetic deformation. Braking is achieved by rotating the twist gear using a small servo motor (“Twist Servo” in Figure 3.1, GoBilda 2000 Series). At approximately 135° of twist, the structure jams against the insert, creating a passive brake. The actuator remains locked until the twist is reversed.

Table 3.1: Robot Electrical and Mechanical Parameters

Thigh Length	$L_{\text{thigh}} = 7\text{cm}$
Shank Length	$L_{\text{shank}} = 15\text{cm}$
Rod Length	$L_{\text{rod}} = 30\text{cm}$
Cart Mass	$m_c = 1.1\text{kg}$
Foot Mass	$m_f = 0.2\text{kg}$
Leg Motor Reflected Inertia	$J = 3.5 \times 10^{-3}\text{kgm}^2$
Leg Motor Gear Ratio	1 : 6
Twist Servo Gear Ratio	1 : 4
Supply Voltage	24V
Motor Constant	$105 \frac{\text{Vs}}{\text{rpm}}$
Resistance	0.143Ω
Peak Rated Power	500 W
Peak Rated Torque	16 Nm

3.4 Hardware and Electronics




Figure 3.4: The HSA-PEA Hopping robot is shown in two poses demonstrating auxetic deformation of the parallel HSA.

The fully assembled robot is presented in Figure 3.4. Two poses are shown to illustrate the spring-like behavior of the HSA; as the leg compresses, the HSA deforms along a twist-constrained auxetic trajectory. In this picture, 400 milligrams of added weights are located above the leg motor and behind the HSA.

The power, computation, and control electronics are shown in Figure 3.4. Power is provided by a benchtop two-quadrant programmable supply set to 24 Volts and software limited to 30 Amps. Two power distribution boards (MJBots PowerDist r4.5b) are used for quiescent power monitoring; one board supplies the Raspberry Pi microcomputer and QDD100 motor controller, the other supplies 7.4 Volts to the twist servomotor through a DC-DC converter. The QDD100 (leg motor) is controlled via CAN by a Raspberry Pi microcomputer running a low-latency Linux kernel at 667Hz. The twist servo is controlled via PWM signals from an Arduino microcontroller, which in turn receives setpoints from the Raspberry Pi via USB. The hopping controller and experimental protocols are implemented as Python programs running on the Raspberry Pi; the motor telemetry and control

commands with the QDD100 are exchanged via the open source Moteus controller API.

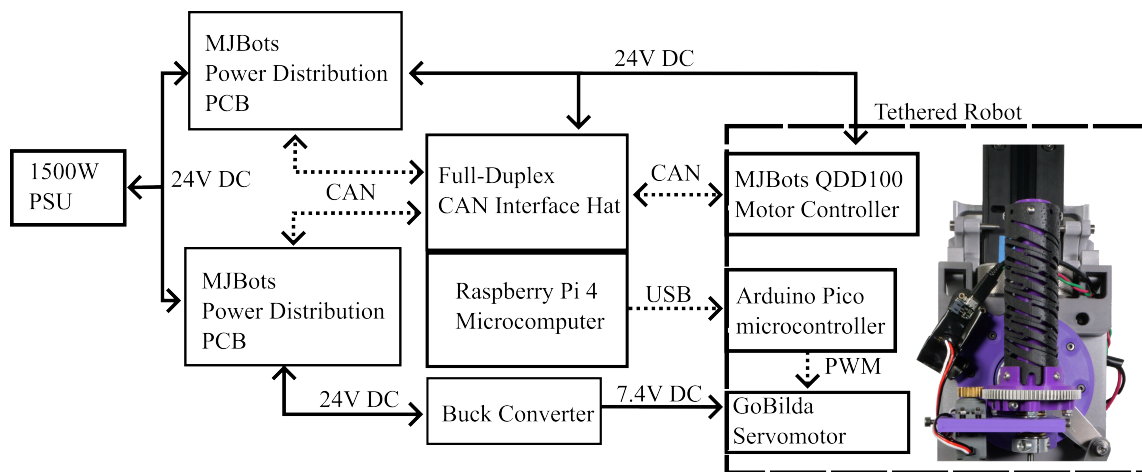


Figure 3.5: Electrical system block diagram. Solid arrows represent power interconnections, dashed arrows denote data and control lines. The actuator is controlled by an off-board microcomputer integrated with a full-duplex controller area network (CAN) interface.

3.5 *Static Braking Tests*

We evaluated power efficiency under static loading by comparing the electrical power consumed by the leg motor versus the twist servo during blocked-force trials. In this test we engaged the braking effect and then incrementally increased the applied torque from the leg motor, while recording the power consumed by the leg motor and the twist servo respectively. As shown in Figure 3.6, leg motor power increases quadratically with force, while the twist motor’s power remains nearly constant. This confirms that the jammed HSA can sustain static loads more efficiently. Extrapolation to higher gear ratios suggests near-zero power consumption with appropriate gearing. As shown in Figure 3.6, the power consumed by the leg motor increases quadratically with blocked force, while the twist motor follows a linear trend. These results show a jammed HSA can support static loads with significantly lower power consumption than the leg motor motor. A projection for a 1:8 gear reduction (solid gold line) illustrates how additional gearing further improves the braking efficiency. Theoretically, using a worm gear a jammed HSA can brake for zero electrical cost.

3.6 *Model-Free Hopping Experiments*

3.6.1 *Methods*

We evaluated the improvement in electrical efficiency provided by the HSA spring through a series of hopping experiments with increasing added mass. These experiments were divided into two groups:

- **With HSA** - The HSA was installed on the actuator (labeled “with HSA” in Figure 3.7).
- **Without HSA** - The leg motor imposed virtual compliance using proportional angle control (labeled “without HSA”).

Both groups utilized a hybrid PD controller to stabilize the hopping behavior, divided into ‘flight’ and ‘stance’ modes. The purpose of the flight control is to lift the foot off the

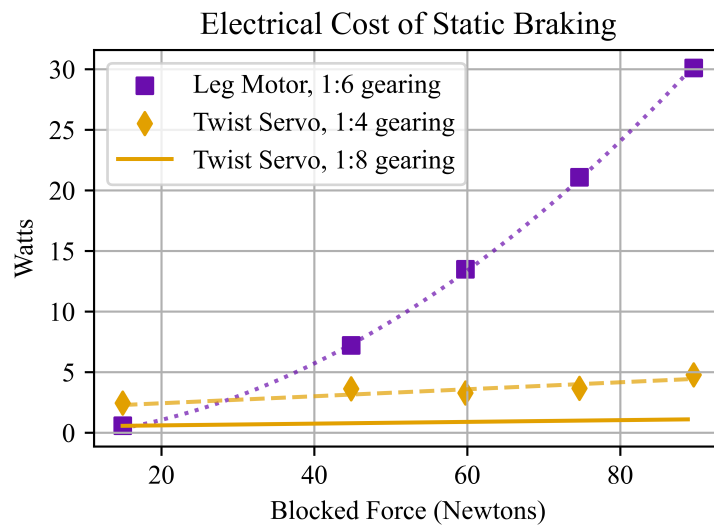


Figure 3.6: Comparison of the electrical power drawn vs. the force blocked by the leg motor compared to the HSA twist motor. The power consumed by the primary motor (squares) is fit by a quadratic function of the blocked force (purple dotted line). The power consumed by the HSA twist motor (diamonds) follows a linear trend (gold dashed line). The solid gold line predicts the power that would be consumed by the twist servo at 1:8 gear reduction.

ground, repositioning for the following stance phase. To achieve the desired hopping height, a fixed push-off torque is applied in the second half of the stance phase.

During the experiments, we collected telemetry from the leg motor, including angle (θ), torque (τ) and current (I), which we used to calculate motor work, joint work and the total electrical losses. In our analysis, we define electrical losses as the sum of net motor work and thermal losses (heat dissipated in the motor windings). Thermal losses in the twist servo were negligible as a result of the high mechanical advantage and friction in its gearbox. The **electrical power** of the leg motor is approximated by

$$P_{\text{elec}} = P_{\text{Therm}} + P_{\text{Motor}} = RI^2 + \tau\dot{\theta}. \quad (3.2)$$

To fairly assess the benefits of the HSA, we included electrical regeneration from negative motor work (i.e., regenerative braking) in our analysis, consistent with modern motor controllers like the QDD100.

3.7 Results

3.7.1 3D Printed HSA-PEA Improves Hopping Efficiency

For each tested mass, between 64 and 77 hops were performed with an HSA and without. The total mean and variance of the hop heights were $5.2 \text{ cm} \pm 2 \text{ mm}$. For each condition, we estimated the mean cost-of-transport with 99% confidence intervals via the bootstrapping method. We found the HSA consistently reduced the hopping cost-of-transport across all tested masses. As stated in section IV, our energy calculation assumes that negative motor work is perfectly regenerated. Without an HSA, the leg motor performs nearly all the negative mechanical work in the system.

3.7.2 Increased Efficiency is Driven by Reduced Joule Heating

Our calculations show an HSA improves the efficiency of hopping by reducing the thermal cost-of-transport (Joule heating), while simultaneously increasing the mechanical cost-of-transport (net motor work).

Figure 3.7(b) shows the reduction in the thermal COT due to Joule heating alongside

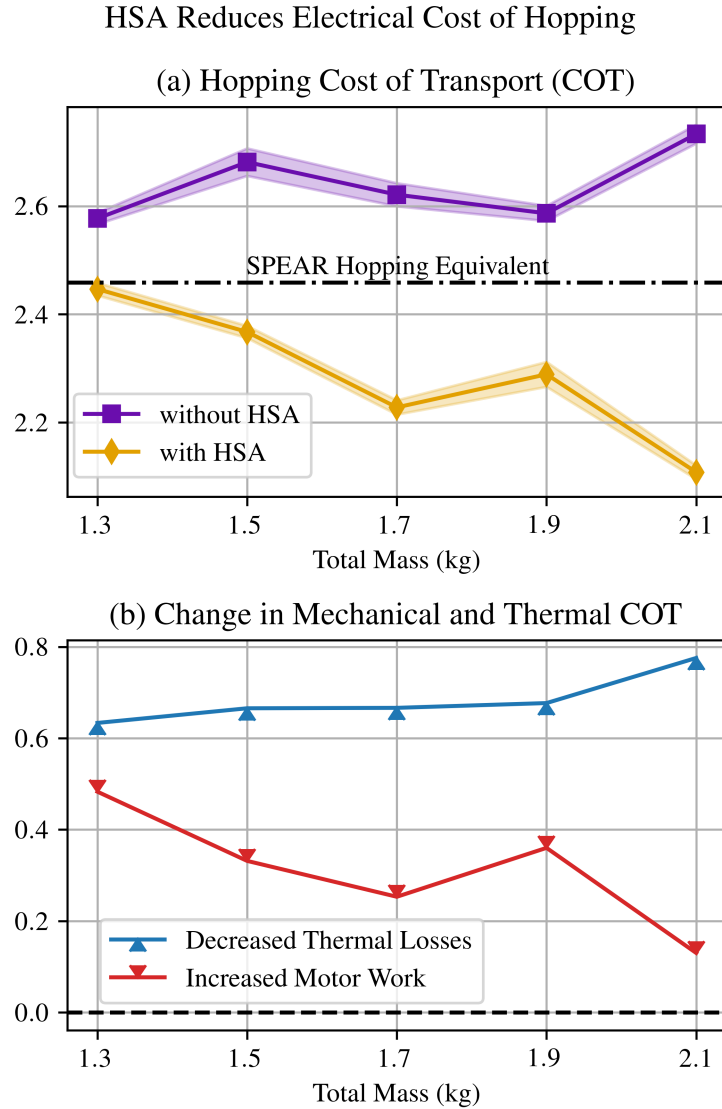


Figure 3.7: Part (a): Hopping cost-of-transport with and without an HSA compared to results from the SPEAR hopping robot. Part (b): Changes to the mechanical and thermal components of COT (with an HSA minus without an HSA), shown in *absolute value*, i.e. the blue curve is the *decrease* in the thermal cost, the red curve is the *increase* in mechanical cost. The reduced hopping COT in (a) is due to reduced thermal cost, while the mechanical cost increases, suggesting energy is dissipated by the HSA.

Table 3.2: SPEAR Vertical Hopping Data [21]. θ_{knee} is the touchdown angle of the knee joint, L_{td} is the length of the virtual leg at touchdown, E is the electrical energy consumed during a hop, Δx is the maximum compression of the virtual leg, and h is the calculated hop height.

θ_{knee} (rad)	L_{td} (cm)	E (J)	Δx (cm)	h (cm)	COT
0.74	60.7	15	10.6	10.5	2.96
1.04	55.7	16.3	21.8	27.5	1.23
1.40	49.5	20	27.8	34.5	1.20

the increase in mechanical COT due to net motor work. These quantities were obtained via numerical integration of the motor and thermal power defined in (3.2) during each hop.

3.7.3 Cost-of-Transport Comparison with SPEAR Hopper

To compare the HSA’s performance with traditional steel springs, we analyzed results from the Switchable Parallel Elastic Actuator Robot (SPEAR) presented by Liu et al [21]. In that paper, the authors provided the electrical energy E consumed, and force-compression curves $F(x)$ of a *virtual* spring between the robot foot and hip. We estimated the hop height h by calculating the potential energy stored in the virtual leg spring, and subtracting work against gravity for a lumped mass of 4.91 kg (SPEAR’s body weight):

$$h \approx -\Delta x + \frac{1}{mg} \int_0^{\Delta x} F(x) dx.$$

Table 3.2 shows that the COT of SPEAR is linearly related to the hop height ($R^2 = 0.93$). Since these hop heights span a large range of 19% to 63% of the virtual leg length, we interpolated the data in Table 3.2 at an intermediate height $h_{\text{eq}} = 15.8$ cm, calculated by rescaling the average hop height from our data by the ratio of SPEAR’s touchdown leg length to our robot’s touchdown leg length. This resulted in an equivalent COT of 2.46, which is shown as a dashed line in Figure 3.7(a).

3.7.4 3D Printed HSA Dissipates Energy

Using the data collected during hop experiments, we investigated the amount of positive work that is supplied by the HSA to the system. The positive work done by the HSA can be inferred from the ratio of positive work done by the leg motor to the total positive work in the system. This is captured in the spring efficiency metric:

$$\eta = 1 - \frac{\int_0^{T_S} \max(0, P_{\text{Motor}}) dt}{\int_0^{T_S} \max(0, P_{\text{Joint}}) dt} \quad (3.3)$$

where T_S is the duration of the stance period in hopping and P_{Joint} is the joint power given by,

$$P_{\text{Joint}} = (J\ddot{\theta}) \cdot \dot{\theta} + (m_{\text{cart}} + m_{\text{weight}}) (\ddot{x} + g) \cdot \dot{x}.$$

The spring efficiency index η was used in [15, 21] to estimate the amount of positive joint work performed by springs in legged robots. We calculated η to be 29% mean \pm 1.7 std%, significantly less than the spring efficiency of 64% reported for SPEAR, suggesting the 3D-printed HSA dissipates more energy than a steel spring.

3.8 Conclusion

This chapter presented the design and experimental evaluation of a hopping robot featuring a 3D-printed Handed Shearing Auxetic (HSA) structure integrated as a parallel elastic and braking element. Static tests demonstrated that when jammed, the HSA effectively supports high static loads with significantly lower electrical power consumption than the primary leg actuator. This highlights its potential as a passive braking mechanism in energy-constrained robotic systems.

In dynamic hopping experiments, the inclusion of the HSA consistently reduced the electrical cost-of-transport across a range of payloads. Efficiency gains were primarily attributed to reduced Joule heating in the leg motor, even as mechanical work slightly increased. Compared to a baseline control strategy using virtual compliance, the HSA-enabled system demonstrated improved energetic performance without sacrificing hopping height or stability. A benchmark comparison with the SPEAR hopping robot further contextualized these

results. At a normalized hop height, our system achieved a comparable or lower cost-of-transport, despite relying on a 3D-printed polymer spring instead of high-performance steel components.

Finally, analysis of the spring efficiency metric revealed that the HSA dissipates a substantial portion of mechanical energy, achieving a mean efficiency of approximately 29%, notably lower than the 64% efficiency reported for SPEAR's steel springs. This suggests that viscoelastic damping is a prominent feature of the HSA's behavior, providing motivation for the modeling and trajectory optimization efforts presented in the following chapter.

Chapter 4

DATA-DRIVEN VISCOELASTIC MODELING

The goal of this chapter is to develop a structured, physics-informed framework for modeling the behavior of viscoelastic actuators based on Handed Shearing Auxetic (HSA) structures. These actuators exhibit strongly nonlinear and configuration-dependent behavior under dynamic loading, complicating their integration into traditional control and planning pipelines.

To address this, we propose a variational modeling framework that represents the elastic and dissipative forces using a pair of convex potential functions. This approach retains the interpretability and structure of Lagrangian mechanics while offering the flexibility needed to capture complex, hysteretic behavior observed in experimental data. We describe the derivation of the model, the selection of convex basis functions, and the regression procedures used to identify parameters from experimental data. The resulting model enables prediction of actuator forces under dynamic inputs, forming the basis for the trajectory optimization and experimental validation described in the next chapter.

In this chapter, we begin with a review of the Lagrange-D'Alembert principle to establish the foundation for modeling non-conservative systems. We then introduce elastic and dissipation potentials and describe how they are used to represent viscoelastic forces in the HSA actuator. Next, we outline the model identification process, including the formulation of a convex optimization problem for fitting potential parameters to experimental data. Finally, we validate the proposed models on both benchtop and hopping trajectories, comparing predictive performance across a hierarchy of nested models.

4.1 Lagrange-D'Alembert Generalized Action Principle

This section reviews the variational principles that form the foundation of our modeling framework. We begin by revisiting the classical formulation of Lagrangian mechanics and

the principle of stationary action. We then introduce the Lagrange-D'Alembert principle, which generalizes this formulation to systems subject to non-conservative forces—such as the viscoelastic forces we aim to model.

The Lagrangian formulation of mechanics describes a system by considering the difference between kinetic and potential energy:

$$L(q, \dot{q}) = T(q, \dot{q}) - V(q) . \quad (4.1)$$

Here, L is the Lagrangian, T denotes the kinetic energy of the system, and V the potential energy. From L , we define the *action* functional:

$$S[q] := \int_{t_0}^{t_1} L(q(t), \dot{q}(t)) dt , \quad (4.2)$$

which assigns a scalar value to a trajectory $q(t)$ over the time interval $[t_0, t_1]$. Hamilton's principle, also known as the principle of stationary action, asserts that physical trajectories extremize this functional: $\delta S = 0$, where

$$\delta S := \left. \frac{d}{d\epsilon} S[q + \epsilon\eta] \right|_{\epsilon=0} = \int_{t_0}^{t_1} \frac{\partial L}{\partial \dot{q}} \dot{\eta} + \frac{\partial L}{\partial q} \eta dt = 0 \quad (4.3)$$

for all admissible variations $\eta(t)$ satisfying $\eta(t_0) = \eta(t_1) = 0$ and $\dot{\eta}(t_0) = \dot{\eta}(t_1) = 0$. Applying integration by parts yields the classical Euler-Lagrange equation:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = 0 . \quad (4.4)$$

This equation encodes Newton's laws in a coordinate-invariant form and serves as the foundation for many physics-informed modeling approaches.

To extend this framework to include non-conservative forces, such as actuation and damping, we appeal to the Lagrange-D'Alembert principle. Suppose $F(q, \dot{q}, t)$ denotes a generalized external force. Its virtual work over a trajectory is defined as:

$$\delta W[q, \eta] := \int_{t_0}^{t_1} F(q(t), \dot{q}(t), t) \cdot \eta(t) dt . \quad (4.5)$$

Incorporating this into the action variation leads to the generalized action principle:

$$\delta S = -\delta W, \quad (4.6)$$

and consequently to the forced Euler-Lagrange equation:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = -F(q, \dot{q}, t). \quad (4.7)$$

In the following sections, we define the unknown forces F in terms of convex elastic and dissipation potentials. These are designed to model the configuration- and velocity-dependent viscoelastic effects observed in HSA actuators.

4.2 Elastic Potentials

In this section, we describe our reduced-order model for elastic energy storage in the handed shearing auxetic (HSA). Rather than modeling stress and strain at the finite-element level, which is computationally expensive and difficult to identify from data, we introduce an abstracted representation via elastic potentials—functions that map generalized coordinates to stored elastic energy. This approach enables efficient integration with trajectory optimization and supports parameter estimation from data.

For a material with ideal elasticity, the energy stored by a reversible (elastic) deformation is the integral of a positive-definite quadratic form of the strain density. However, modeling the strain density at the level of a finite-element discretization is prohibitively costly in our application. Instead, we adopt a simplified model of elastic energy, analogous to the reduced stiffness models employed by Sharma and Kramer [35]. That is, we seek an *elastic potential* function $P_E : \mathbb{R}^n \rightarrow \mathbb{R}$ with suitably low dimension n . The essential properties of an elastic potential $P(q)$ are:

1. $P(q) > 0$ for all q .
2. $P(q)$ is strongly convex.

The most prominent example of an elastic potential is Hooke's law, which describes the energy stored in a coiled spring with the potential $P(q) = \frac{1}{2}(q - q_0)^\top K(q - q_0)$. The spring

force $F(q)$ is the differential of P , i.e., $F(q) = \frac{\partial P}{\partial q} = K(q - q_0)$. Strong convexity of P implies that the spring constant satisfies $\frac{\partial^2 P}{\partial q^2} = K \succ 0$. In Hooke’s law, the spring force is a linear function of the displacement $\Delta q = q - q_0$. However, for more complex structures, the force resulting from an elastic potential P can be nonlinear; this occurs when the second derivative $\frac{\partial^2 P}{\partial q^2}$ is not constant.

4.3 Dissipation Potentials

We now introduce dissipation potentials—functions that model energy loss due to velocity-dependent friction. This complements the elastic potentials described above and allows us to capture viscoelastic behavior using a unified variational framework.

Our prior work suggests that 3D printed handed shearing auxetics are dissipative. We aim to quantify this dissipation using a velocity-dependent model. Velocity-dependent friction is a common source of dissipation in mechanical systems. Such forces have been integrated with Lagrangian mechanics since at least the 1870s [28], when Lord Rayleigh introduced a concept now known as a *dissipation potential* [24]. For a mechanical system with configuration q of dimension n , Rayleigh’s original formulation was:

$$R(\dot{q}) = \frac{1}{2} \sum_{i=1}^n c_{ij} \dot{q}_i \dot{q}_j = \frac{1}{2} \dot{q}^\top C \dot{q}, \quad (4.8)$$

where the matrix $C \succeq 0$. While the elastic potentials we previously examined measure the *energy* stored by an elastic deformation, the Rayleigh potential measures the *power* dissipated due to friction. Since their introduction, Rayleigh potentials have been extended to model many kinds of friction [24]. We consider a Rayleigh potential $R(q, \dot{q})$ to have three essential properties:

1. $R(q, \dot{q}) \geq 0$ for all (q, \dot{q}) .
2. $R(q, \dot{q})$ is strongly convex in \dot{q} .
3. $R(q, 0) = 0 < R(q, \dot{q})$ for all $\dot{q} \neq 0$.

In analytical mechanics, the generalized force produced by a dissipation potential $R(q, \dot{q})$ is prescribed under the Lagrange-D'Alembert principle as:

$$F_d(q, t) = \left. \frac{\partial R}{\partial \dot{q}} \right|_{q(t), \dot{q}(t)}. \quad (4.9)$$

Consequently, for a system with Lagrangian $L(q, \dot{q})$, the modified Euler-Lagrange equations are:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = - \frac{\partial R}{\partial \dot{q}}. \quad (4.10)$$

4.4 Viscoelastic HSA Model

Our model of the HSA consists of an elastic potential $P(q, \theta)$ and a dissipation potential $R(q, \dot{q}, \theta)$, where $q \in \mathbb{R}$ is the length of the HSA, \dot{q} is the length rate-of-change, and $\theta \in \mathbb{R}$ is the twist angle that parametrizes the auxetic trajectory. These potentials are constructed from linear combinations of basis functions:

$$\begin{aligned} P(q, \theta) &= x_1 q + x_2 q^2 - \sum_{i=1}^n x_{i+2} e^{-\rho_i(q, \theta)} \\ &= \sum_{i=1}^{n+2} x_i \phi_i(q, \theta), \quad x \succeq 0 \end{aligned} \quad (4.11)$$

$$\begin{aligned} R(q, \dot{q}, \theta) &= y_1 \dot{q}^2 + \sum_{i=1}^n y_{i+1} \dot{q}^2 e^{-\rho_i(q, \theta)} \\ &= \sum_{i=1}^{n+1} y_i \psi_i(q, \dot{q}, \theta), \quad y \succeq 0 \end{aligned} \quad (4.12)$$

Observe that Hooke's law for the length parameter q is spanned by $\{\phi_1 = q, \phi_2 = q^2\}$, and the classical Rayleigh potential is spanned by $\psi_1 = \dot{q}^2$ (assuming $\dot{\theta} = 0$, see footnote¹). The $e^{-\rho_i}$ are Gaussian kernels with the distance metrics:

$$\rho_i(q, \theta) := \frac{(q - \mu_{q,i})^2}{2\sigma_{q,i}^2} + \frac{(\theta - \mu_{\theta,i})^2}{2\sigma_{\theta,i}^2}$$

¹The twist angle of the HSA θ is fixed in every trajectory we have tested, so we do not treat it as a dynamic variable. This limitation could be removed by introducing new basis functions $\{\theta, \theta^2, \dot{\theta}^2, \dot{\theta}^2 e^{-\rho_i}\}$ into the viscoelastic model.

The hyperparameters $(\mu_{q,i}, \mu_{\theta,i})$ determine the kernel centers, and $(\sigma_{q,i}, \sigma_{\theta,i})$ control their width. Notably, a non-positive sum of Gaussians is not convex. However, convexity can be enforced on a compact domain $\mathcal{D} := [q_{\min}, q_{\max}] \times [\theta_{\min}, \theta_{\max}]$ by choosing $\sigma_{q,i} = \sup_{q \in \mathcal{D}} |q - \mu_{q,i}|$, and likewise for $\sigma_{\theta,i}$. By relaxing these constraints, the same basis can be used to model inelastic potential functions.

The models we present in the results section use ten Gaussian kernels equally spaced on a 5×2 grid spanning $q_i \in [0.1 \text{ m}, 0.18 \text{ m}]$ and $\theta_i \in \{0, 67.5^\circ \times \pi/180^\circ\}$. For all kernels, we use

$$\begin{aligned}\sigma_q &= \frac{2\Delta q}{\sqrt{3}} = \frac{2}{\sqrt{3}} \times \frac{0.18\text{m} - 0.1\text{m}}{4} \approx 0.23\text{m} \\ \sigma_\theta &= \frac{2\Delta\theta}{\sqrt{3}} = \frac{2}{\sqrt{3}} \frac{\pi \times 67.5^\circ}{180^\circ} \approx 1.36 \text{ radians.}\end{aligned}$$

These parameters allow the curvature of P to vary while keeping its minimum curvature close to x_1 , which encourages (but does not enforce) strong convexity in P through the constraint $x \succeq 0$. In particular, the maximum curvature of each Gaussian in the q dimension occurs at $q^* = \mu_{q,i} \pm \sqrt{3}\sigma_q = \mu_{q,i} \pm 2\Delta q$, which is coincident with the minimum curvature of another Gaussian or lies outside the domain $[0.1\text{m}, 0.18\text{m}]^2$. The same holds for $\mu_{\theta,i}$ in relation to σ_θ .

4.5 Data-Driven Model Identification

This section presents our approach for identifying the parameters of a viscoelastic HSA model from actuator data. Given observed trajectories of length and motor force at a fixed twist angle, we derive a cost function that quantifies how well a candidate model satisfies the forced Euler-Lagrange equations. The model parameters are identified by solving a convex optimization problem, formulated as a constrained, weighted least-squares fit.

We assume an observed trajectory $q : [a, b] \rightarrow \mathbb{R}$ and an applied motor force $u : [a, b] \rightarrow \mathbb{R}$, for some fixed twist angle θ . In the following, we use a semicolon to separate the arguments of higher-order functions into *parameters* and *signals* or *functions*. This notation defines

²We suspect that the minimum curvature of P in this method is bounded by a linear function of x specific to the grid spacing, so that global convexity may be enforced without compromising problem (4.19) in III.B.

functionals in a way that distinguishes between data and parameters. As an example, suppose that $q(t)$ is a continuous path in \mathbb{R} . The expression $f(\tau; q) := \int_a^\tau q(t)dt$ defines a higher-order function: each τ gives a functional $f(\tau; \cdot)$, and each path $q(t)$ gives a function $f(\cdot; q)$.

We begin with a partial Lagrangian $L(q, \dot{q})$ that describes the rigid-body mechanics of the actuator, excluding viscoelastic effects. Using the notation introduced above, we expand the Euler-Lagrange equation into an equivalent time-dependent form:

$$E_L(t; q) := \left(\frac{\partial^2 L}{\partial \dot{q}^2} \Big|_{q(t), \dot{q}(t)} \ddot{q}(t) \right) + \left(\frac{\partial^2 L}{\partial q \partial \dot{q}} \Big|_{q(t), \dot{q}(t)} \dot{q}(t) \right) - \left(\frac{\partial L}{\partial q} \Big|_{q(t), \dot{q}(t)} \right) = 0. \quad (4.13)$$

We now include the applied motor force $u(t)$ and the viscoelastic model, represented by elastic and dissipation potentials $P = \sum x_i \phi_i$ and $R = \sum y_i \psi_i$, as generalized forces. Let $\Phi(t; q, \theta) \in \mathbb{R}^{n+2}$ and $\Psi(t; q, \theta) \in \mathbb{R}^{n+1}$ compute the generalized force corresponding to each basis function of P and R :

$$\Phi_j(t; q, \theta) := \frac{\partial \phi_j}{\partial q} \Big|_{q(t), \theta} \quad (4.14)$$

$$\Psi_k(t; q, \theta) := \frac{\partial \psi_k}{\partial \dot{q}} \Big|_{q(t), \dot{q}(t), \theta} \quad (4.15)$$

The error in the forced Euler-Lagrange equation is then:

$$e(x, y, t; q, \theta, u) := \Phi(t; q, \theta) \cdot x + \Psi(t; q, \theta) \cdot y + u(t) - E_L(t; q). \quad (4.16)$$

A perfect model of the system would satisfy the generalized action principle $\delta S + \delta W = 0$ for all variations, implying that the residual $e(x, y, t; \cdot)$ is zero for all t . To evaluate the accuracy of a candidate model with weights (x, y) , we compute the L_2 norm of the residual:

$$\int_a^b \|e(x, y, t; q, \theta, u)\|^2 dt \quad (4.17)$$

So far, we have assumed that $q(t)$ is a continuous trajectory, allowing derivatives $\dot{q}(t)$ and $\ddot{q}(t)$ to be computed directly. In practice, only discrete samples of $q(t)$ and $u(t)$ are available. We fit continuously differentiable polynomial splines to these data, which smooth noise and regularize the derivatives. Let $\{(q_i, \theta_i, u_i)\}_{i=1}^N$ be a collection of spline segments, each a polynomial of degree d on domain (t_{i-1}, t_i) . For each i , we approximate the integral using Gauss-Legendre quadrature:

$$\int_a^b \|e(x, y, t; q, \theta, u)\|^2 dt \approx \sum_{j=1}^{2d} w_j \|A_{ij}x + B_{ij}y - c_{ij}\|^2$$

Here, (A_{ij}, B_{ij}, c_{ij}) evaluate $(\Phi, \Psi, u - E_L)$ at the j -th Legendre root τ_j on interval (t_{i-1}, t_i) with corresponding weight w_j :

$$\begin{aligned} A_{ij} &:= \Phi(\tau_j; q_i, \theta_i) \\ B_{ij} &:= \Psi(\tau_j; q_i, \theta_i) \\ c_{ij} &:= u_i(\tau_j) - E_L(\tau_j; q_i) \end{aligned}$$

The total cost over the full dataset is:

$$\sum_{i=1}^N \sum_{j=1}^{d_i} w_j \|A_{ij}x + B_{ij}y - c_{ij}\|^2 . \quad (4.18)$$

To estimate the model parameters, we solve the optimization problem:

$$\begin{aligned} \min_{x, y} \quad & \sum_{i=1}^N \sum_{j=1}^{d_i} w_j \|A_{ij}x + B_{ij}y - c_{ij}\|^2 \\ \text{such that} \quad & x \succeq 0 \in \mathbb{R}^{n+2} \\ & y \succeq \in \mathbb{R}^{n+1} \end{aligned} \quad (4.19)$$

Problem (4.19) is a constrained and weighted least-squares regression, which we solve using CVXPY in Python with default settings.

4.6 Viscoelastic Model Validation

The goal of our model validation experiments is to determine whether these generalized, nonlinear behaviors are supported by experimental data. We consider four models arranged in a nested hierarchy:

1. (**HE+RD**) is a minimal model of **H**ookean **E**lasticity and classical **R**ayleigh **D**issipation,
2. (**HE+GD**) features **H**ookean **E**lasticity and **G**eneralized **D**issipation,
3. (**GE+RD**) features **G**eneralized **E**lasticity and classical **R**ayleigh **D**issipation,
4. (**GE+GD**) the full model of **G**eneralized **E**lasticity and **D**issipation.

These relationships form a nested structure: 1) HE+RD is a strict subspace of both Models 2) HE+GD and 3) GE+RD, which are themselves subspaces of 4) GE+GD.

To evaluate model quality, we first fit each model to a set of 290 hopping trajectories and perform likelihood ratio tests on a held-out set of 130 trajectories. This allows us to determine whether the added complexity of nonlinear components significantly improves prediction. Specifically, we test the following nested pairs: (HE+RD vs. HE+GD), (HE+GD vs. GE+GD), (HE+RD vs. GE+RD), and (GE+RD vs. GE+GD), isolating the effects nonlinear elasticity and varying damping coefficients respectively.

Because Models 2–4 include more parameters than the baseline (HE+RD), we also test for overfitting by training each model on a separate benchtop dataset collected under idealized conditions. We then evaluate generalization by comparing their predictive performance on hopping data. The benchtop dataset was collected by suspending the robot leg in air and commanding the actuator through a sinusoidal trajectory using the onboard motor. This procedure was repeated for each HSA twist angle.

4.7 Results

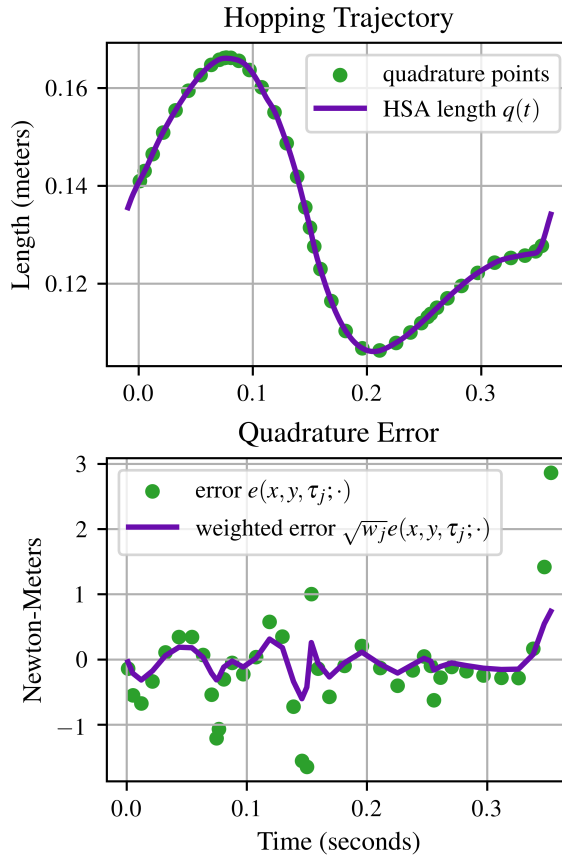


Figure 4.1: Illustration of the integral quadrature method for evaluating model error. Top: Measured HSA length over a single hop (purple) with quadrature points (green). Bottom: Residuals of the Euler-Lagrange equation (green) and weighted residuals (purple). A systematic error appears at landing due to missing ground impulse data but has reduced influence due to low quadrature weight.

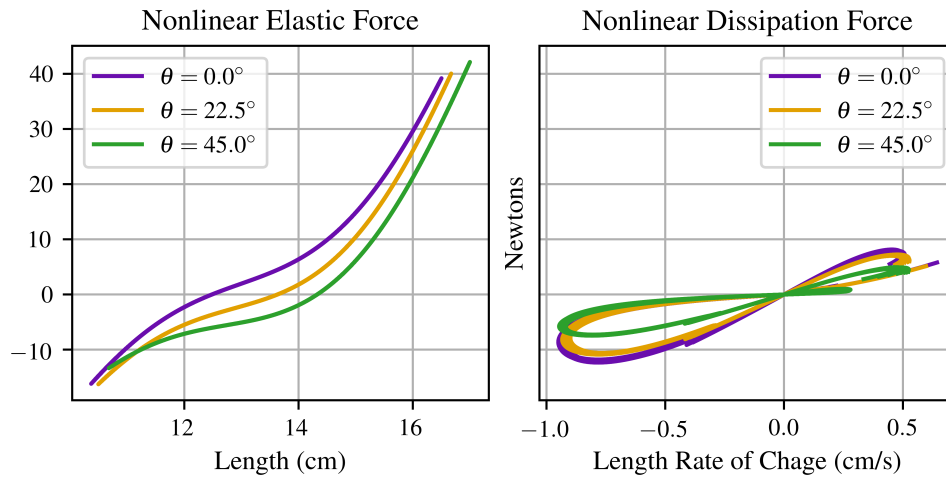


Figure 4.2: Model-predicted elastic (left) and dissipative (right) forces across all test hops. Each trace represents a different HSA twist angle θ . Nonlinear behavior is evident in both force types: elastic force shows varying curvature, and damping force exhibits configuration and rate dependent hysteresis.

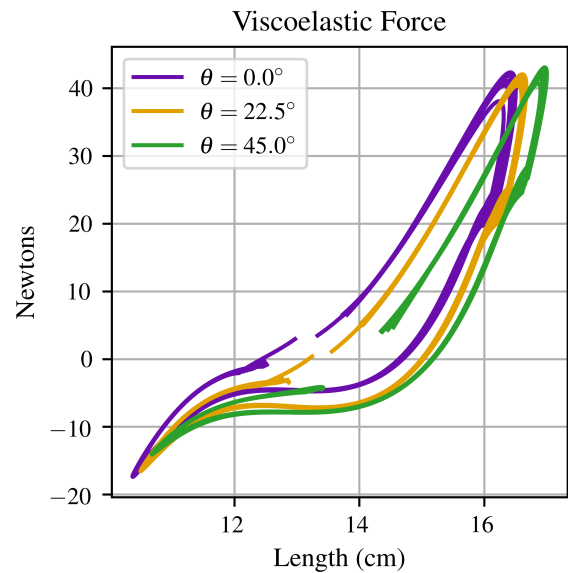


Figure 4.3: Model-predicted viscoelastic force across all test hops. Each trace shows the sum of the contributions from the elastic and dissipation potential at a different HSA twist angle θ .

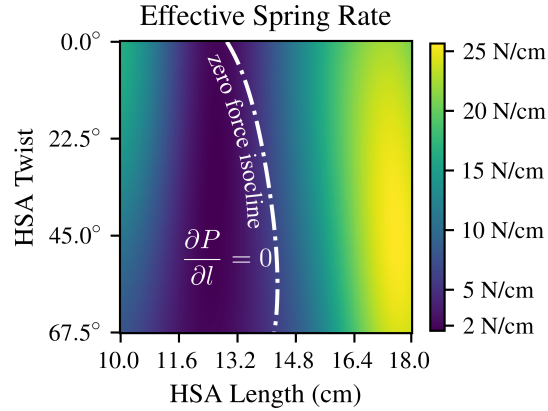


Figure 4.4: Heatmap of the effective spring rate for the elastic potential in model (1). The dashed line shows the zero-force length at each twist angle. The spring rate is strictly positive in accordance with strong convexity.

Table 4.1: Comparison of prediction residuals for different HSA models. The meaning of the model labels (i.e. HE+RD) is outlined in section 4.6. The left sub-table reports results from models trained on hopping data and tested on holdout hopping data. The right sub-table reports results of models trained on benchtop data and tested on holdout hopping data.

Model	Trained on Hopping Data				Trained on Benchtop Data			
	RMSE	R ²	Kurtosis	Skew	RMSE	R ²	Kurtosis	Skew
HE+RD	28	0.90	2.9	1.6	28	0.91	3.8	1.8
HE+GD	26	0.92	2.7	1.6	28	0.91	3.9	1.9
GE+RD	19	0.96	0.5	0.6	22	0.94	2.3	1.5
GE+GD	16	0.97	0.6	0.4	21	0.95	2.1	1.4

Table 4.2: Summary of hypothesis tests comparing nested HSA models. Tests compare reduced models (left) against more expressive alternatives (right). Permutation tests are used where residuals deviate significantly from normality; otherwise, likelihood ratio tests (LRT) are applied. For each test, the null hypothesis assumes no improvement in fit. Critical values correspond to the $\alpha = 0.05$ significance level.

Test	Type	Null Hypothesis	Statistic	Critical Value
HE+RD vs. HE+GD	Permutation	HE+GD does not improve fit	0*	-49
HE+GD vs. GE+RD	Permutation	GE+RD does not improve fit	0*	-49
GE+GD vs. GE+RD	Likelihood Ratio	GE+GD does not improve fit	47.3*	3.94

4.8 Discussion

The model validation results further support the physical realism and utility of the proposed framework. As shown in Figure 4.2 (left), the learned elastic potential predicts a continuously varying elastic curvature along the auxetic trajectory and correctly anticipates zero-force length increases with twist. Figure 4.2 (right) shows that the effective damping coefficient of the dissipation potential varies continuously along auxetic trajectories.

The total force predicted by the model is presented in Figure 4.3, which demonstrates that the model captures hysteretic effects typical of viscoelastic materials. The effective spring rate (i.e., the curvature of the elastic potential) and predicted equilibrium length are presented in Figure 4.4, showing that the learned potential is strongly convex.

Quantitative model validation results are summarized in Table 4.1 and Table 4.2. Nested model comparisons are presented in Table 4.2, with all tests showing statistical significance. Model error (RMSE) is measured in terms of change in action (milli-Joule-seconds). For a typical hop lasting 400 milliseconds, one unit of this error corresponds to an average of 2.5 Newton-meters of unexplained torque. These results confirm that the modeling framework accurately reproduces actuator behavior. Across model comparisons, the most substantial improvements in prediction accuracy arose from using a nonlinear elastic potential with varying curvature—a direct analogue of a nonlinear generalization of Hooke’s law. This improvement persisted even when models were trained on benchtop data and tested on hopping trajectories, demonstrating robustness to distribution shift.

By contrast, nonlinear dissipation kernels contributed less consistently to predictive performance. The fully generalized model GE+GD outperformed GE+RD in all conditions. However, HE+GD did not outperform HE+RD when trained on benchtop data. These results suggest that nonlinear dissipation effects may capture subtle dynamics only when paired with a nonlinear elastic force model.

4.9 Limitations

A primary limitation of our current modeling approach is the treatment of the HSA twist as a static parameter rather than a dynamic configuration variable. This simplification

was imposed by a hardware constraint in our platform: the low-bandwidth servomotor used to control twist angle. To fully characterize the coupled dynamics of the HSA, future work should treat twist and length on equal footing, potentially using a higher-bandwidth actuator capable of closed-loop position and torque control.

A related limitation is our use of scalar Hookean elasticity and classical Rayleigh dissipation. A more complete treatment would include coupling terms such as $q\theta$ and $\dot{q}\dot{\theta}$, as well as quadratic terms like θ^2 and $\dot{\theta}^2$. This is equivalent to defining Hookean elasticity and Rayleigh dissipation using symmetric positive-definite matrices $K \succ 0$ and $C \succ 0$, as in the work of Sharma et al. [35]. The cross-coupling term $q\theta$, in particular, may help explain the variation in zero-force length at different twist angles [8, 37].

Our basis function representation for both the elastic and dissipation potentials was fixed throughout training. While sufficient for the current analysis, future improvements could involve hyperparameter optimization—such as learning kernel locations and widths—or expanding the function class to include richer families of convex functions.

Moreover, the viscoelastic models were trained entirely using motor telemetry. While this indirect measurement captures global actuator behavior, it introduces significant noise into velocity and acceleration estimates, which may bias the predicted force. Direct sensing of strain may improve model identifiability and reduce reliance on overparameterization. Finally, while residual errors are lowest during mid-stance, they remain elevated at impact and takeoff—likely due to missing ground reaction force data. More accurate or redundant sensing during these phases could enhance both training and validation performance.

4.10 Variational Interpretation and Future Extensions

The variational structure underlying our modeling framework not only supports principled identification but also enables deeper insight into model behavior and error structure. In this subsection, we revisit the cost function from the perspective of the action principle, highlighting a saddle-point formulation in which the model parameters and the path variation interact adversarially.

Recall the definition of the loss function in equation (4.18):

$$\sum_{i=1}^N \sum_{j=1}^{d_i} w_j \|A_{ij}x + B_{ij}y - c_{ij}\|^2 .$$

Earlier in this chapter we showed that this corresponds to the loss function of a weighted least-squares regression. However, if we interpret this loss in terms of the variation of action $\delta S + \delta W$, we see that the loss above is given by a particular spline of path variations $\eta_i(t)$:

$$\begin{aligned} & \sum_{j=1}^{2d} w_j \|A_{ij}x + B_{ij}y - c_{ij}\|^2 \\ &= \sum_{j=1}^{2d} w_j \langle \eta_i(\tau_j), A_{ij}x + B_{ij}y - c_{ij} \rangle \\ &\approx \int_{t_{i-1}}^{t_i} \langle \eta_i(t), e(x, y, t; q_i, \theta_i, u_i) \rangle dt \end{aligned}$$

Here, $\eta_i(t)$ is the unique degree- $2d+4$ polynomial satisfying the boundary and interpolation constraints:

$$\begin{aligned} \eta_i(t_{i-1}) &= \eta_i(t_i) = 0 , \\ \dot{\eta}_i(t_{i-1}) &= \dot{\eta}_i(t_i) = 0 , \\ \eta_i(\tau_j) &= e(x, y, \tau_j; q_i, \theta_i, u_i) = A_{ij}x + B_{ij}y - c_{ij} \end{aligned}$$

This construction highlights the connection between problem (4.19) and the generalized action principle: problem (4.19) is equivalent to minimizing the change in action $\delta S + \delta W$ along a path variation η determined by the model structure. It also suggests a more robust alternative: minimizing the maximum change in action $\delta S + \delta W = \langle \eta(t), e(x, y, t; \cdot) \rangle_2$ over

bounded variations η . This leads to a saddle-point problem with convex variables (x, y) and concave variable η .

The saddle-point formulation of the identification problem, where the path variation $\eta(t)$ acts as a concave adversarial variable, provides a valuable new lens for understanding model performance. In this view, $\eta(t)$ highlights regions of the trajectory where model error contributes most significantly to the change in action. This role is analogous to influence functions in robust statistics or weighting kernels in locally weighted regression. In particular, the following saddle problem closely resembles a least-squares regression that is robust to the sample weights [32]:

$$\min_{x,y} \max_{\eta(t)} \int_{t_0}^{t_1} \langle \eta(t), e(x, y, t; q, \theta, u) \rangle dt \quad (4.20)$$

The shape of $\eta(t)$ may carry useful diagnostic information. For instance, a spatially biased η may reveal model misspecification or systematic underfitting in certain regions of the configuration space. Likewise, local oscillations in η can indicate that the current quadrature rule or polynomial resolution is insufficient to resolve the model dynamics in that region, suggesting a direction for adaptive refinement.

These insights point to several promising directions for future work, including adaptive quadrature schemes informed by $\eta(t)$, or regularized saddle formulations that improve robustness to noise and model mismatch. They also reinforce the physical interpretability of the variational framework: the learned model not only predicts force and energy, but also shapes the geometry of permissible path variations through the action landscape.

Chapter 5

OPTIMIZED HOPPING EXPERIMENTS

This chapter presents a case study in model-based trajectory optimization for a legged robot driven by a viscoelastic actuator. Building on the identification and validation results of Chapter 3, we apply the learned HSA model to synthesize energy-efficient hopping behaviors. By leveraging the structure-preserving formulation of the actuator dynamics, we generate cyclic motions that minimize thermal losses and exploit passive compliance.

We begin by formulating a trajectory optimization problem grounded in the learned viscoelastic model, incorporating actuator and task constraints. Optimized trajectories are then deployed in hardware and compared to open-loop hops to assess gains in performance and efficiency. Finally, we discuss practical implications, limitations, and the observed trade-offs between energy savings and control fidelity in the presence of actuator hysteresis. These results illustrate the benefits of integrating physically meaningful, data-driven models into the planning loop for dynamic legged locomotion.

5.1 Problem Formulation

The optimized hopping trajectories are computed by solving a direct collocation problem. This problem jointly determines the trajectory $q(t)$ and applied motor force $u(t)$ during the *stance* period that minimizes total electrical energy consumption. For simplicity, the flight phase is not modeled because it does not contribute significantly to total motor work or electrical loss.

In this model, electrical power is the sum of mechanical power, $u(t)\dot{q}(t)$, and thermal power $J(u(t), q(t))$. The integral of mechanical power determines the net electrical energy converted into useful work, accounting for electrical regeneration due to negative motor work. The integral of thermal power gives the energy lost to Joule heating, which is strictly

positive. The problem definition is:

$$\min_{q(t), u(t)} \int_0^{t_f} u \dot{q} + J(u, q) dt \quad (5.1)$$

subject to the constraints:

$$q(0) = q(t_f) \quad (\text{boundary conditions})$$

$$\dot{q}(0) = -\dot{q}(t_f) = v$$

$$0 = e(x, y, t; q, \theta, u) \quad (\text{dynamics})$$

$$\tau_{\min} \leq \tau(u(t), q(t)) \leq \tau_{\max} \quad (\text{torque limits})$$

We solve this problem by representing the trajectory and applied force as third-degree polynomial splines with one knot at $t_f/2$. Continuity in q is enforced up to the second derivative, while no continuity constraints are imposed on u . The cost function is evaluated using Gauss-Legendre quadrature, and the constraints are enforced at the roots of fourth-degree Legendre polynomials on $[0, t_f/2]$ and $[t_f/2, t_f]$. We use the SLSQP method in SciPy to solve this problem (cite SciPy and Kraft’s paper).

The boundary condition $\dot{q}(0) = -\dot{q}(t_f) = v$ determines the hopping height. In contrast, the HSA twist θ and the terminal position $q(0) = q(t_f)$ are not optimized; they are chosen experimentally for gait stability. This is necessary because we use a state-event driven control law: the feedforward force $\nu(t)$ is applied when the state $q(t)$ crosses q_0 with positive velocity. However, since the rest length of the HSA varies significantly with twist θ , an improperly chosen θ can trigger stance control while the foot is still airborne, resulting in unstable behavior that diverges from the reference trajectory.

5.2 Data Collection and Statistical Analysis

The primary goal of the hopping experiments was to quantify the effect of the HSA structure on hopping efficiency, as measured by the Cost of Transport (COT). In previous work, we demonstrated that the HSA significantly reduced energy expenditure during hopping. However, those experiments used a fixed, non-optimized control policy, leaving open the question of whether HSA-induced improvements would persist under optimal actuation.

To address this, we repeated the hopping experiments from our prior study using trajectories optimized for energy efficiency. The resulting dataset includes four experimental conditions: unoptimized control with an HSA, denoted $(\text{HSA}, \overline{\text{OPT}})$; unoptimized control without an HSA, $(\overline{\text{HSA}}, \overline{\text{OPT}})$; optimized control with an HSA, (HSA, OPT) ; and optimized control without an HSA, $(\overline{\text{HSA}}, \text{OPT})$. The unoptimized control law was taken directly from our previous design study (citation), while the optimized trajectories were obtained by solving the minimum-energy control problem defined in (5.1).

For each condition, we collected a minimum of 60 hops at each of five payload masses: 1.3, 1.5, 1.7, 1.9, and 2.1 kilograms. In the optimized HSA condition, we further varied the HSA structure by testing three twist angles: 0, 22.5, and 45 degrees. For each hop, we computed the hop height h from the observed velocity of the leg at liftoff, and calculated the total electrical energy consumption E from motor telemetry. The Cost of Transport was then computed as $\text{COT} := \frac{E}{mgh}$, where E matches the cost function of problem (5.1).

To evaluate the relative contributions of the HSA and trajectory optimization to hopping efficiency, we conducted a regression analysis across the four experimental groups. However, the raw data presented several statistical challenges. The recorded hops were neither independent nor identically distributed, due to sequential dependencies, HSA hysteresis, variation across nominally identical HSA devices, and minor differences in controller parameters. Furthermore, the distributions of payload mass and hop height were not balanced across conditions, and both are known to influence COT.

To address these issues, we used a bootstrap resampling strategy to compare distributions of sample means across experimental conditions. We employed importance sampling to equalize the marginal distribution of payload masses across groups and validated this balancing using a chi-squared contingency test. A linear regression was then performed on the bootstrapped sample means, with covariates for mean hopping height, HSA condition, optimization status, and their interaction effects.

5.3 Results

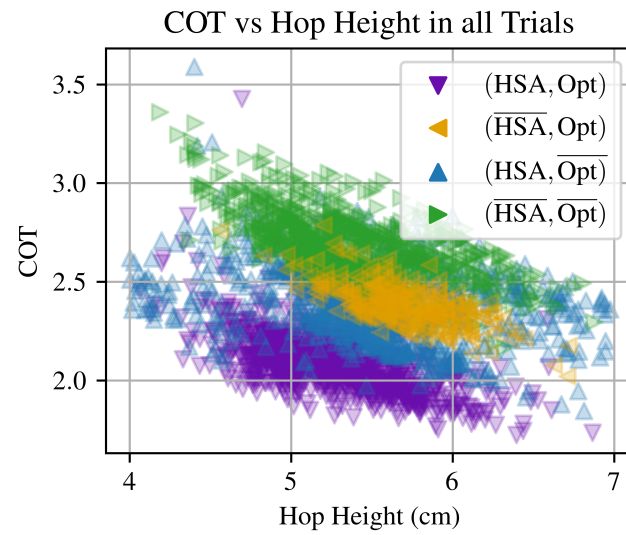


Figure 5.1: Scatter plot of Cost of Transport versus hopping height for all trials. Each condition is color-coded and partially transparent to visualize density. Overlap across groups and high variance motivate the need for regression to isolate effects of optimization and the HSA.

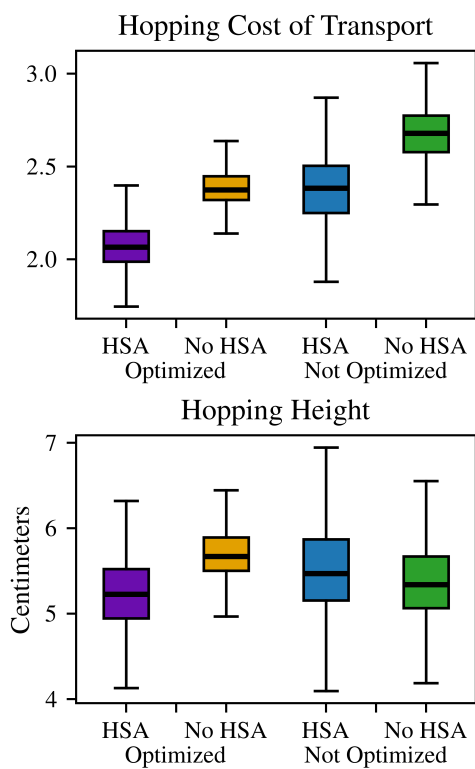


Figure 5.2: Boxplots comparing Cost of Transport and hopping height across the four experimental conditions: HSA vs. no HSA and optimized vs. not optimized. While both optimization and the HSA reduce mean COT, large and unequal variances complicate direct comparisons. Differences in hopping height across groups motivate controlling for this covariate in regression analysis.

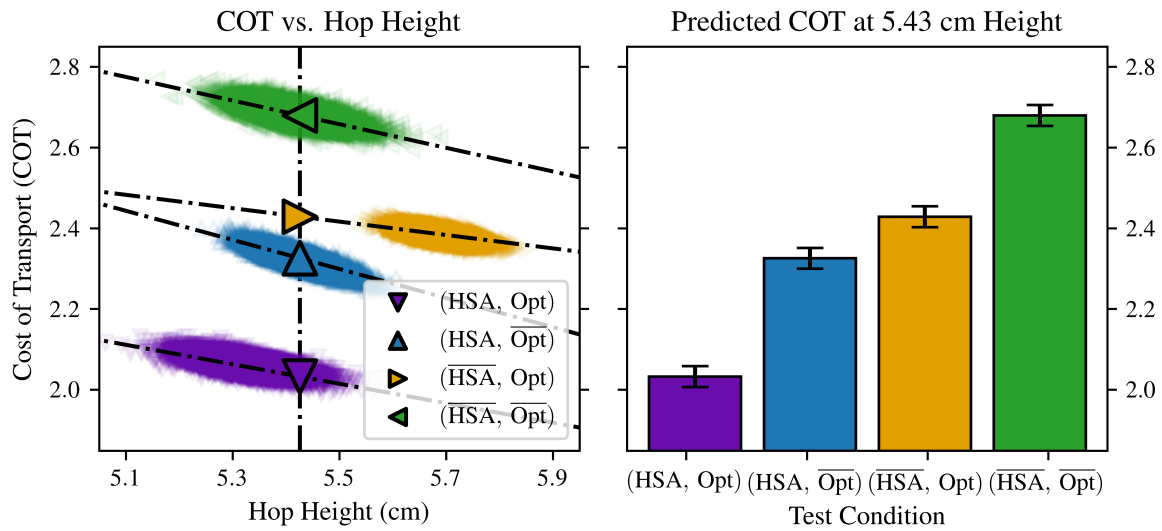


Figure 5.3: Left: Bootstrapped sample means ($n = 40$, 10,000 resamples) of COT versus hopping height for each experimental group. Regression lines highlight trends. Right: Predicted COT at the median hopping height (5.43 cm) with 95% confidence intervals. Both optimization and the HSA significantly reduce COT, and the benefit of the HSA is amplified under optimized trajectories.

Table 5.1: Linear regression results predicting mean Cost of Transport (COT) from hopping height, HSA condition, and optimization status over 40,000 resamples. The model achieves $R^2 = 0.996$. Residuals are approximately normal (excess kurtosis = 0.20, skew = 0.55).

Predictor	Coef.	<i>t</i> -value	95% CI
HSA	0.013	0.89	[-0.015, 0.041]
Opt	-0.940	-66.1	[-0.968, -0.912]
Mean Height h	-29.3	-146	[-29.7, -28.9]
$h \times$ HSA	-6.76	-25.7	[-7.27, -6.24]
$h \times$ Opt \times $\overline{\text{HSA}}$	12.7	50.1	[12.2, 13.2]
$h \times$ Opt \times HSA	11.92	45.2	[11.4, 12.4]

5.4 Discussion

In Chapter 2, we demonstrated that incorporating a Handed Shearing Auxetic (HSA) structure into the leg actuator significantly reduced the Cost of Transport (COT) during unoptimized hopping. This improvement was attributed primarily to reduced Joule heating in the primary leg motor. However, mechanical work calculations in that study indicated that the HSA also dissipated a substantial amount of energy, raising concerns that this benefit might diminish under optimized control strategies.

Our present results clarify that picture. While the proposed viscoelastic HSA model confirms the presence of damping effects via a non-conservative dissipation potential, the regression analysis of hopping experiments revealed that the energetic benefit of the HSA actually increased when paired with optimized control. This finding suggests that allowing for slightly higher net motor work, in exchange for reduced resistive losses, is a favorable trade-off. It reinforces the value of incorporating parallel elasticity in quasi-direct drive leg actuators, where energy efficiency hinges on balancing control effort with compliant energy storage.

Chapter 6

CONCLUSION

This dissertation presented a comprehensive study on the design, viscoelastic modeling, and control of a novel hopping robot incorporating a Handed Shearing Auxetic (HSA) structure. Our experimental proof of concept demonstrated how 3D-printed architected materials can be integrated with quasi-direct drive actuators to provide both mechanical compliance and braking functionality. Through analysis of the robot’s hopping Cost of Transport (COT), we showed that HSAs can serve as an effective alternative to conventional springs. This opens a promising direction for research and innovation at the intersection of compliant manipulation and legged robotics, motivating further research into the trade-offs and performance limits of 3D-printed architected materials.

To further evaluate the utility of the proposed actuator, we developed a system identification method based solely on motor telemetry, which revealed the actuator’s nonlinear viscoelastic properties. These models were then integrated into a trajectory optimization pipeline to enhance the energy efficiency of hopping. Taken together, this work offers a holistic framework that supports rapid prototyping, experimentation, characterization, and optimization of novel elastic actuators incorporating complex viscoelastic structures.

6.1 Summary of Contributions

We began by developing a novel actuator that placed a Handed Shearing Auxetic (HSA) structure in parallel with a quasi-direct drive motor (HSA-PEA), and leveraged the auxetic trajectory to lock the joint. Building on this, we constructed a hopping robot within a controlled testbed. As presented in Chapter 2, this study demonstrated that the elastic compliance provided by the HSA reduced the Cost of Transport (COT) during hopping. We also showed that the auxetic trajectory can be exploited to drastically reduce static power consumption under heavy loads. Energy analysis revealed that the improved dynamic

efficiency resulted primarily from reduced Joule heating in the primary motor. However, we also observed an increase in total motor work, suggesting that the HSA itself dissipated a significant amount of energy—an insight that directly motivated our subsequent work on model identification and trajectory optimization.

Chapter 3 introduced a novel physics-informed identification method grounded in the variational principles of analytical mechanics. Building upon prior work, our method is distinguished by a convex formulation that minimizes action variation over non-invasive elastic energy and dissipation functions. By casting model identification as a least-squares regression on the Euler-Lagrange residual, and expressing elastic and dissipative forces as differentials of convex potentials, we obtained interpretable and physically consistent models that generalized across operating conditions. We validated these models using both benchtop and locomotion data, confirming their robustness under distribution shift and their ability to reproduce key features of actuator behavior, including hysteresis and strain-dependent stiffness.

In Chapter 4, we leveraged these models to explore the dynamic efficiency of the HSA-PEA through trajectory optimization. The resulting hopping behaviors were deployed on hardware and compared across four conditions: with and without HSAs, and with optimized and non-optimized control strategies. Our statistical analysis confirmed that both trajectory optimization and HSA structures independently reduced energy consumption, with a synergistic effect when combined. These findings validate the practical utility of our modeling framework and provide further evidence that HSAs represent a promising direction for innovation in leg actuator design.

6.2 Limitations

Several limitations of this work should be acknowledged. First, further investigation is needed into the fatigue behavior and longevity of 3D-printed HSAs under the cyclic loading characteristic of legged locomotion. Second, while our modeling framework successfully captures nonlinear elasticity and viscoelastic damping, it treats HSA twist as a static parameter rather than a dynamic configuration variable. This simplification—necessitated by hardware constraints—limits our ability to fully characterize the coupled dynamics of the

HSA.

Third, the optimization framework used a fixed collocation mesh and did not explicitly model the flight phase of hopping, potentially introducing suboptimality near stance-transition boundaries. Fourth, model identification relied solely on motor telemetry, which introduces noise in velocity and acceleration estimates and may bias the inferred force dynamics. Finally, while our convex basis function representation enabled efficient and interpretable modeling, it did not account for bilinear coupling between length and twist—a potentially important feature of HSA mechanics.

6.3 Future Directions

This work opens several avenues for future research. A natural extension is to treat twist as a dynamic configuration variable, enabling joint optimization of axial and torsional motions. Realizing this would require higher-bandwidth actuation and sensing, as well as extensions to both the model structure and optimization framework. Second, adaptive mesh refinement and quadrature schemes could improve numerical fidelity in both trajectory optimization and model fitting, particularly by leveraging the residual-weighting insights introduced in Chapter 3.

Third, combining indirect identification via motor telemetry with direct strain sensing could reduce reliance on high-order polynomial fitting and improve robustness to measurement noise. The variational framework also invites integration with tools from robust control and estimation. For instance, saddle-point formulations or minimax strategies could be used to support robust regression in the presence of model mismatch or unmodeled dynamics.

Finally, the structured modeling approach developed here could be extended to a wider class of architected materials and hybrid actuators, enabling systematic design and control of next-generation robotic limbs that integrate compliance, damping, and geometric tunability in a unified mechanical structure.

6.4 Closing Remarks

Taken together, the results of this dissertation demonstrate that physically informed, data-driven modeling can bridge the gap between material nonlinearity and robotic control. By aligning the structure of learning algorithms with the variational principles of mechanics, we can capture complex actuator behaviors in a form that is amenable to optimization, interpretation, and real-world deployment. As robotics increasingly embraces architected structures and heterogeneous materials, such approaches may offer a blueprint for the principled integration of mechanics, data, and design.

BIBLIOGRAPHY

- [1] Alejandro F. Azocar, Luke M. Mooney, Levi J. Hargrove, and Elliott J. Rouse. Design and characterization of an open-source robotic leg prosthesis. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 111–118. ISSN: 2155-1782.
- [2] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *113(15):3932–3937*.
- [3] Samuel A. Burden, Thomas Libby, Kaushik Jayaram, Simon Sponberg, and J. Maxwell Donelan. Why animals can outrun robots. *9(89):eadi9754*. Publisher: American Association for the Advancement of Science.
- [4] Zixi Chen, Federico Renda, Alexia Le Gall, Lorenzo Mocellin, Matteo Bernabei, Théo Dangel, Gastone Ciuti, Matteo Cianchetti, and Cesare Stefanini. Data-driven methods applied to soft robot modeling and control: A review. *22:2241–2256*.
- [5] Lillian Chin, Michelle C. Yuen, Jeffrey Lipton, Luis H. Trueba, Rebecca Kramer-Bottiglio, and Daniela Rus. A simple electric soft robotic gripper with high-deformation haptic feedback. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2765–2771. ISSN: 2577-087X.
- [6] Corentin Coulais, Eial Teomy, Koen de Reus, Yair Shokef, and Martin van Hecke. Combinatorial design of textured mechanical metamaterials. *535(7613):529–532*. Publisher: Nature Publishing Group.
- [7] 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. ISSN: 2153-0866.
- [8] Ian Good, Tosh Brown-Moore, Aditya Patil, Daniel Revier, and Jeffrey Ian Lipton. Expanding the design space for electrically-driven soft robots through handed shearing auxetics. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10951–10957.
- [9] Sam Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks.

- [10] Duncan W. Haldane, M. M. Plecnik, J. K. Yim, and R. S. Fearing. Robotic vertical jumping agility via series-elastic power modulation. 1(1):eaag2048. Publisher: American Association for the Advancement of Science.
- [11] Mokarram Hossain, Rukshan Navaratne, and Djordje Perić. 3d printed elastomeric polyurethane: Viscoelastic experimental characterizations and constitutive modelling with nonlinear viscosity functions. 126:103546.
- [12] Christian Hubicki, Jesse Grimes, Mikhail Jones, Daniel Renjewski, Alexander Spröwitz, Andy Abate, and Jonathan Hurst. ATRIAS: Design and validation of a tether-free 3d-capable spring-mass bipedal robot. 35(12):1497–1521.
- [13] Jonathan W. Hurst. The role and implementation of compliance in legged locomotion. ISBN: 9798617016828.
- [14] Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C. Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, Remo Diethelm, Samuel Bachmann, Amir Melzer, and Mark Hoepflinger. ANYmal - a highly mobile and dynamic quadrupedal robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 38–44. ISSN: 2153-0866.
- [15] Marco Hutter, C. David Remy, Mark A. Hoepflinger, and Roland Siegwart. Efficient and versatile locomotion with highly compliant legs. 18(2):449–458. Conference Name: IEEE/ASME Transactions on Mechatronics.
- [16] Pranav Kaarthik, Francesco L. Sanchez, James Avtges, and Ryan L. Truby. Motorized, untethered soft robots via 3d printed auxetics. 18(43):8229–8237. Publisher: Royal Society of Chemistry.
- [17] Navvab Kashiri, Andy Abate, Sabrina J. Abram, Alin Albu-Schaffer, Patrick J. Clary, Monica Daley, Salman Faraji, Raphael Furnemont, Manolo Garabini, Hartmut Geyer, Alena M. Grabowski, Jonathan Hurst, Jorn Malzahn, Glenn Mathijssen, David Remy, Wesley Roozing, Mohammad Shahbazi, Surabhi N. Simha, Jae-Bok Song, Nils Smit-Anseeuw, Stefano Stramigioli, Bram Vanderborght, Yevgeniy Yesilevskiy, and Nikos Tsagarakis. An overview on principles for energy efficient robot locomotion. 5. Publisher: Frontiers.
- [18] Jeffrey Ian Lipton, Robert MacCurdy, Zachary Manchester, Lillian Chin, Daniel Cellucci, and Daniela Rus. Handedness in shearing auxetics creates rigid and compliant structures. 360(6389):632–635. Publisher: American Association for the Advancement of Science.
- [19] Siyu Liu, Jiatao Ding, Chunlei Lu, Zhirui Wang, Bo Su, and Zhao Guo. A novel optimization design of dual-slide parallel elastic actuator for legged robots. 29(4):2886–2894.

- [20] Xin Liu and Ioannis Poulakakis. On the energetics of a switchable parallel elastic actuator design for monopedal running. In *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 769–774. IEEE.
- [21] Xin Liu, Anthony Rossi, and Ioannis Poulakakis. A switchable parallel elastic actuator and its application to leg design for running robots. *23(6):2681–2692*.
- [22] Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning.
- [23] Lucas R. Meza, Alex J. Zelhofer, Nigel Clarke, Arturo J. Mateos, Dennis M. Kochmann, and Julia R. Greer. Resilient 3d hierarchical architected metamaterials. *112(37):11502–11507*. Publisher: Proceedings of the National Academy of Sciences.
- [24] E Minguzzi. Rayleigh’s dissipation function at work. *36(3):035014*. Publisher: IOP Publishing.
- [25] Fernández Pelayo, David Blanco, Pedro Fernández, Javier González, and Natalia Beltrán. Viscoelastic behaviour of flexible thermoplastic polyurethane additively manufactured parts: Influence of inner-structure design factors. *13(14):2365*. Number: 14 Publisher: Multidisciplinary Digital Publishing Institute.
- [26] Michiel Plooi, Glenn Mathijssen, Pierre Cherelle, Dirk Lefeber, and Bram Vanderborght. Lock your robot: A review of locking devices in robotics. *22(1):106–117*. Conference Name: IEEE Robotics & Automation Magazine.
- [27] Michiel Plooi, Wouter Wolfslag, and Martijn Wisse. Clutched elastic actuators. *22(2):739–750*. Conference Name: IEEE/ASME Transactions on Mechatronics.
- [28] John William Strutt (3rd Baron Rayleigh.) and John William Strutt Rayleigh. *The Theory of Sound*. Macmillan. Google-Books-ID: pPICD9D7tKcC.
- [29] Jacob Reher, Wen-Loong Ma, and Aaron D. Ames. Dynamic walking with compliance on a cassie bipedal robot. In *2019 18th European Control Conference (ECC)*, pages 2589–2595.
- [30] Federico Renda, Costanza Armanini, Vincent Lebastard, Fabien Candelier, and Frederic Boyer. A geometric variable-strain approach for static modeling of soft manipulators with tendon and fluidic actuation. *5(3):4006–4013*.
- [31] M. C. Saldívar, E. Tay, A. Isaakidou, V. Moosabeiki, L. E. Fratila-Apachitei, E. L. Doubrovski, M. J. Mirzaali, and A. A. Zadpoor. Bioinspired rational design of bi-material 3d printed soft-hard interfaces. *14(1):7919*. Publisher: Nature Publishing Group.

- [32] Philipp Schiele, Eric Luxenberg, and Stephen Boyd. Disciplined saddle programming.
- [33] Sangok Seok, Albert Wang, Meng Yee Chuah, David Otten, Jeffrey Lang, and Sangbae Kim. Design principles for highly efficient quadrupeds and implementation on the MIT cheetah robot. In *2013 IEEE International Conference on Robotics and Automation*, pages 3307–3312. ISSN: 1050-4729.
- [34] Maziar Ahmad Sharbafi, Mohammad Javad Yazdanpanah, Majid Nili Ahmadabadi, and Andre Seyfarth. Parallel compliance design for increasing robustness and efficiency in legged locomotion—proof of concept. *24(4):1541–1552*. Conference Name: IEEE/ASME Transactions on Mechatronics.
- [35] Harsh Sharma and Boris Kramer. Preserving lagrangian structure in data-driven reduced-order modeling of large-scale dynamical systems. *462:134128*.
- [36] Brian M. de Silva, Kathleen Champion, Markus Quade, Jean-Christophe Loiseau, J. Nathan Kutz, and Steven L. Brunton. PySINDy: A python package for the sparse identification of nonlinear dynamics from data.
- [37] Maximilian Stölzle, Lillian Chin, Ryan L. Truby, Daniela Rus, and Cosimo Della Santina. Modelling handed shearing auxetics: Selective piecewise constant strain kinematics and dynamic simulation. In *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 1–8. ISSN: 2769-4534.
- [38] Maximilian Stölzle, Daniela Rus, and Cosimo Della Santina. An experimental study of model-based control for planar handed shearing auxetics robots. volume 30, pages 153–167.
- [39] Ryan L. Truby, Lillian Chin, and Daniela Rus. A recipe for electrically-driven soft robots via 3d printed handed shearing auxetics. *6(2):795–802*. Conference Name: IEEE Robotics and Automation Letters.
- [40] Xin Liu, Anthony Rossi, and Ioannis Poulakakis. SPEAR: A monopodal robot with switchable parallel elastic actuation. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5142–5147. IEEE.
- [41] Yevgeniy Yesilevskiy, Weitao Xi, and C. David Remy. A comparison of series and parallel elasticity in a monopod hopper. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1036–1041. ISSN: 1050-4729.
- [42] Xiaoyu Zheng, Howon Lee, Todd H. Weisgraber, Maxim Shusteff, Joshua DeOtte, Eric B. Duoss, Joshua D. Kuntz, Monika M. Biener, Qi Ge, Julie A. Jackson, Sergei O. Kucheyev, Nicholas X. Fang, and Christopher M. Spadaccini. Ultralight, ultrastiff mechanical metamaterials. *344(6190):1373–1377*. Publisher: American Association for the Advancement of Science.

- [43] Xiaoyu Zheng, William Smith, Julie Jackson, Bryan Moran, Huachen Cui, Da Chen, Jianchao Ye, Nicholas Fang, Nicholas Rodriguez, Todd Weisgraber, and Christopher M. Spadaccini. Multiscale metallic metamaterials. 15(10):1100–1106. Publisher: Nature Publishing Group.