

MODELING AND GAIT CONTROL FOR PRINCIPALLY KINEMATIC  
LOCOMOTION SYSTEMS

by  
Siming Deng

A dissertation submitted to The Johns Hopkins University in conformity  
with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland  
January, 2024

© 2024 Siming Deng  
All rights reserved

# Abstract

Geometric mechanics provides a framework through which top-down insights permit novel motion planning approaches to a broad class of locomotion systems with symmetry, including analytical computation of optimal gaits. A core premise of this framework is that complex locomotor mechanics can be rewritten in a kinematic form, where momentum effects are negligible (e.g., dissipative forces dominate the physical interaction between the body and the environment). Here, we consider a subclass of such group-invariant dynamical systems; the equations of motion can be kinematically reduced such that the body velocity is expressed as a shape-dependent linear mapping of shape velocity. The same framework has been instrumental in understanding cyclic locomotion where a local model can be constructed in the neighborhood of the observed limit cycle, using data points from stochastically perturbed, repeated behaviors. We aim to tackle practical challenges in implementing data-driven geometric methods, including (1) enhancing the framework’s real-time capabilities, (2) further improving data efficiency, (3) providing a principled gait modulation scheme for continuous steering control, and (4) extending the framework to even more general systems.

We have developed an adaptive system identification extension to the current framework that enables the real-time update of the system model that can be used to modify behaviors iteratively within a behavior optimization scheme. This capability not only enhances fundamental behaviors but also enables precise motion tracking. Additionally, we showcase its usage in refining behaviors, aiding in injury recovery, and adapting to different terrains, especially scenarios where simulations (models built

from first principles) provide inadequate guidance for real-world situations.

Obtaining a locomotion model efficiently from data, as described above, is an important step that can be leveraged to guide robot behavior, as it facilitates the principled designing of control policies. Acquiring nominal gaits for various behavioral goals is an important step; in practical scenarios, the ability to steer from these nominal gaits is just as important. We introduce principled gait modulation algorithms designed to modify a nominal gait for single-parameter steering control by constructing a continuum gait library using either global or local model information. This approach opens doors to motion planning and control for systems with complex, *a priori* unknown dynamics in which an intuitive ‘joystick’ control can be provided.

As soft (and in general underactuated) robots gain popularity, efficiently acquiring useful reduced-order models for these highly underactuated systems becomes increasingly demanding. Geometric methods have great potential in describing the behaviors of many systems within this category. We developed an extension to the data-driven geometric framework that builds layers of interconnected models, which include actuator and locomotor dynamics obtained from data gathered during repeatedly stochastically perturbed behaviors. These linked models are grounded in the general formulations of Lagrangian systems with symmetry, making them suitable for a broad spectrum of robots with first-order, low-pass actuator dynamics, such as hydrogel crawlers powered by swelling-based actuators. These models effectively encapsulate the dynamics of system shape and body movements in a simplified swimming robot model. We also suggested a numerical optimization scheme for control signals via iterative model refinement, which we employed to optimize the input waveform for the hydrogel crawler.

# Thesis Readers

Dr. Noah Cowan (Primary Advisor)

Professor

Department of Mechanical Engineering

Johns Hopkins University

Dr. Louis Whitcomb

Professor

Department of Mechanical Engineering

Johns Hopkins University

Dr. Brian Bittner

Robotics Research Scientist

Johns Hopkins University Applied Physics Lab

*To my family*

# Acknowledgements

I would like to express my sincere gratitude to all those who have supported and witnessed my journey at Hopkins. First and foremost, I am deeply thankful to my advisor, Noah Cowan, for his guidance, encouragement, and support throughout the course of my study. He has admiring expertise across disciplines and insightful perspectives on both engineering and science. During the initial Covid outbreak, he was very supportive in helping me to stay on track with my research and life. To me, he has been both a mentor and a friend. I would like to thank Brian Bittner, whose knowledge and insights have provided a solid foundation for my research. I am also grateful to Louis Whitcomb for his valuable advice and feedback on this dissertation.

I feel fortunate to be part of the LIMBS Lab and to have the opportunity to work with such a great group of people that become my academic family. Thank you to the LIMBS Lab present and past members: Manu Madhav, Ravi Jayakumar, Ismail Uyanik, Debojyoti Biswas, Gorkem Secer, Ted Oh, Shahin Lashkari, Yu Yang, Di Cao, Michael Wilkinson, Brian Woronowicz, and Murtaza Hathiyari. Despite the fact that our workplaces are scattered all over the campus, I enjoyed the time we spent together sharing ideas and discussing research.

I am grateful for collaborations with Vicky Nguyen, Rebecca Schulman, David Gracias, and Ross Hatton, whose knowledge and perspectives have broadened my horizons. I appreciate the opportunity to see the world over the shoulders of giants. I would like to acknowledge my colleagues and fellow researchers from the EFRI group, Bibekananda Datta, Aishwarya Pantula, Ruohong Shi, Junning Liu, and Charlie Chen.

It was fun and enlightening to work with you even though we have different academic backgrounds. Thanks as well to Marin Kobilarov and Chen Li, for their valuable feedback on my first GBO exam. Thanks to Simon Lenard and Anna Goodridge for their help with the software and hardware issues. I learned a lot from you. Thanks to the administrative staff at LCSR and the Mechanical Engineering Department, especially to Alison Morrow, Jordan Card, and Mike Bernard, for answering my administrative questions and helping me with placing orders. Thank you to all my other great friends at Hopkins. I will remember those trails we explored, those concerts we enjoyed, those restaurants we discovered, and those soccer games we watched. I look forward to our paths crossing again in the future.

Finally, I want to thank my family for their unconditional love and support. Their belief in my abilities has been a constant source of motivation, and their support has sustained me during challenging times. To my parents, I am so lucky to have you as my first teacher and my longest-standing friend in life.

# Contents

<b>Abstract</b> . . . . .	<b>ii</b>
<b>Dedication</b> . . . . .	<b>v</b>
<b>Acknowledgements</b> . . . . .	<b>vi</b>
<b>Contents</b> . . . . .	<b>viii</b>
<b>List of Figures</b> . . . . .	<b>xi</b>
<b>Chapter 1 Introduction</b> . . . . .	<b>1</b>
1.1 Bio-inspired Robotic Locomotion . . . . .	2
1.2 Modeling Locomoting Systems . . . . .	5
1.2.1 Bottom-up: Induction from First Principles . . . . .	6
1.2.1.1 Geometric Mechanics . . . . .	6
1.2.2 Top-down: Data-driven Methods . . . . .	9
1.2.3 On the Spectrum of White-box and Black-box . . . . .	10
1.2.3.1 Data-driven Geometric Methods . . . . .	11
1.3 Steering and Control of Nonholonomic Systems . . . . .	13
1.4 Contributions . . . . .	15
1.5 Dissemination . . . . .	16
<b>Chapter 2 In-situ System Identification for Principally Kinematic Systems</b> . . . . .	<b>18</b>



2.1	Motivation . . . . .	18
2.2	Background . . . . .	21
2.2.1	Principally Kinematic Systems . . . . .	21
2.2.2	Data-driven methods for geometric systems . . . . .	23
2.3	Methods . . . . .	24
2.3.1	Adaptive Filters . . . . .	24
2.3.2	Model Metrics . . . . .	26
2.3.3	Online Noise Generation . . . . .	27
2.4	Accuracy of Adaptive Geometric Modeling . . . . .	28
2.5	Adaptation to Changing Environments . . . . .	30
2.6	In-situ Gait Optimization with Adaptive Geometric Modeling . . . . .	33
2.7	Discussion and Conclusion . . . . .	35
<b>Chapter 3 Enhancing Maneuverability via Gait Design . . . . .</b>		<b>40</b>
3.1	Motivation . . . . .	40
3.2	Background . . . . .	42
3.2.1	Locomotion Model . . . . .	42
3.2.2	Gait Optimization . . . . .	44
3.3	Maneuverability . . . . .	46
3.3.1	Gradient-Based Approaches . . . . .	46
3.3.2	Pareto Front . . . . .	47
3.3.3	Online Control . . . . .	48
3.4	Results . . . . .	49
3.5	Demonstration . . . . .	52
3.6	Discussion and Conclusion . . . . .	52
<b>Chapter 4 Data-Driven Geometric Modeling of Systems with Low-Bandwidth Actuators . . . . .</b>		<b>55</b>

4.1	Motivation . . . . .	55
4.2	Background . . . . .	59
4.2.1	Geometric Locomotion Model . . . . .	59
4.2.2	Data-driven Modeling . . . . .	60
4.3	Methods . . . . .	62
4.3.1	Low-bandwidth Shape Control . . . . .	62
4.3.2	Optimizing Behaviors and Iterative Model Refinement . . . . .	64
4.4	Illustrative Example: Purcell Swimmer with Low-bandwidth Actuation . . . . .	65
4.4.1	Input Generation . . . . .	66
4.5	Main Application: Hydrogel Crawler . . . . .	67
4.5.1	Actuator Dynamics . . . . .	67
4.5.2	Hydrogel Crawler . . . . .	68
4.5.3	Finite Element Model . . . . .	68
4.5.4	Input Generation . . . . .	70
4.5.5	Shape Parameterization . . . . .	71
4.5.6	Input Optimization . . . . .	72
4.6	Discussion and Conclusion . . . . .	74
	<b>Chapter 5 Discussion and Future Work . . . . .</b>	<b>77</b>
	<b>References . . . . .</b>	<b>80</b>

# List of Figures

**Figure 2-1** The Purcell swimmer incrementally experiences more stochastically perturbed cycles in a new operating environment. Model metrics of the batch model (green) and adaptive model (red) are evaluated for their predictive performance at that point in the experiment on a hold-out test set of 40 cycles (for 100 pairs of training and testing trials). The model metrics for both methods are shown in boxplots where each box corresponds to 5, 25, 50, 75, 95 percentiles of a model trained with the first  $n \in [5, 10, 15, 20, 25, 30, 35, 40]$  cycles of data experienced. . . . 28

**Figure 2-2** The predictive quality of the adaptive models (blue,  $\lambda_{RLS} = 0.99$ , and orange, ( $\lambda_{RLS} = 0.7$ )) and batch model (green) are compared throughout environmental changes experienced by the Purcell swimmer. Each model is trained with 40 cycles of data collected in ‘Env 1’. In this test, the system is unknowingly exposed to two other environments, where each environment consists of 40 cycles of data. We changed the drag coefficient ratio experienced on the links for the environment changes, where this ratio is 2.0, 3.0, 4.0, respectively for environments 1, 2, and 3. The batch model (which is never recomputed during the trial) has constant large predictive errors on new substrates. The adaptive model is refined each time a sample is collected and can adapt to the new environments relatively quickly. We note that the forgetful model has a faster adaptation rate but achieves lower, less stable prediction quality throughout the trial. . . . . 31

**Figure 2-3** The predictive quality of the adaptive models (blue,  $\lambda_{RLS} = 0.99$ , and orange, ( $\lambda_{RLS} = 0.7$ )), batch model (green), and ground truth model before injury (red) are compared throughout an injury (links broke off) experienced by the Purcell swimmer. Each model is trained with (a. 50, b. 100) cycles of data collected in normal conditions. In this test, the system originally consists of (a. 4, b. 9) links, and it undergoes an injury, where (a. one, b. two) links break off. The batch model is not updated within the trial and it has constant significant predictive errors in the new conditions. The adaptive model is refined each time a sample is collected and can adapt to the injury relatively quickly for both low and high dimensional systems. The ground truth model is the simulation model that generates the data used for this demonstration, hence the log error of this model before the injury is negative infinity. It quickly degrades after the injury happens. . . . . 32

**Figure 2-4** Above we lay out an optimization architecture for behavior refinement. (Top Level) The optimization will attempt to improve behaviors only when it receives a high enough prediction confidence from the adaptive model. (Middle Level) The adaptive model recursively updates from real-world interactions, also producing an estimate of its model confidence in real-time (quantified as  $\Gamma$ ). (Bottom Level) A nominal gait, parameterized as  $p$ , is given to the *Command Gen* module, which generates stochastic perturbations around the nominal gait. . . . . 33

**Figure 2-5** An illustration of rebasing the local linear model at the nominal gait change can further improve data efficiencies for (a) 3-link and (b) 5-link Purcell Swimmer in gait iterations. For no-rebase trials (black), we let the model adapt its weights to reflect nominal gait changes. For rebase trials (red), the model weights are updated using the previous model predictions on the new nominal gait as described in (2.8). The rebase method is able to achieve similar objective function values with fewer empirical data. 36

**Figure 2-6** Gait optimization results for Purcell swimmers using the adaptive model optimizer. (a) The three-link Purcell swimmer can refine its gaits from grey to black, optimizing in 40 cycles of experience. 50 trials of the (b) three-link, (c) five-link, and (d) nine-link Purcell swimmer optimization for extremality, with current progress plotted throughout the experience of the trial. Relative performance in displacement per cycle is shown compared to the initial behavior, with [5,50,95] percentiles of performance plotted throughout and [5,25,50,75,95] percentile boxplots provided for the end of the run. The width of the boxplot indicates the number of cycles executed before reaching the final performance value. . . . . 37

**Figure 3-1** An overview of an example system and its optimal gaits. (a) Schematic of a three-link viscous swimmer—a minimal template [117] for understanding locomotion. The shape of the system, described by its two joint positions, lives in a two-dimensional space. The coordinate of the body frame is chosen according to the *minimal perturbation* principle [57]. (b) The optimal gait (red) in the sense of  $x$ -direction motion efficiency (i.e. the largest displacement in  $x$ -direction per unit power dissipation). The gait is plotted over the system’s *constraint curvature* height function corresponding to  $x$ -direction movement  $(-D\mathbf{A})^x$ , which provides an approximation to the system’s locomotor behavior. (c) Similar to (b), the optimal gait (blue) for turning motion. The two optimal gaits are cross-illustrated in both (b) and (c) with dashed curves. (d) A trivial demonstration of the system’s spatial trajectory from a gait-switching execution. The black arrows over the trajectory denote the system’s  $x$ -axis at the beginning of each cycle, and the color-coded trajectory represents which gait the system is operating on. . . . . 42

**Figure 3-2** An illustration of gait modulations generated with (a) gradient iteration and (b) first-order approximation at the nominal gait. When the modulation variable  $u = 0$ , the gait (red) is the  $x$ -direction optimal gait (i.e., the same as the one shown in Fig. 3-1(b)). As  $u$  increase in value, the nominal gait is modulated toward the top-right corner for additional  $\theta$  movement. . . . . 49

**Figure 3-3** An overview of the performance of the gait libraries from the global approach. (a) Gaits generated with the global approach, where  $\alpha = 1, 0$  corresponds to those optimal gaits shown in Fig. 3-1(b-c). (b) Displacement in  $(x, y, \theta)$  per cycle of the gait library. Note that the discontinuity appears around the region  $\alpha = 0.4$  where the optimizer decides to qualitatively change the gait. (c) Pareto front, generated from (3.13), on the efficiency trade-off map. The discontinuity appears here as a large gap where the optimizer was not able to reliably converge to data points in between, likely due to a bifurcation in the objective function as a function of  $\alpha$ . . . . . 50

**Figure 3-4** An efficiency comparison between the gait modulations generated by local and global approaches. The performance of the gradient-based methods is almost as good as the global optimizer locally at the x-optimum, and slowly drops off. . . . . 51

**Figure 3-5** A demonstration of gait modulation in motion. (a) Shape trajectory, as quasi-periodic functions, of the robot over a normalized time (i.e. one full lap in the shape space per normalized time). (b) The system's spatial trajectory by executing the above shapes. Black arrows denotes the system's x-axis position at the beginning of each lap, and the color coded trajectory shows approximately (not to scale) the average turning curvature is currently being performed. The shape space trajectory of each lap is also shown above the trajectory as a simple visualization. 54



<b>Figure 4-1</b>	Experimental screenshots of thermal cycling of two-segment robots with a flexible linker at the end of a cooling half-cycle (top), and heating half-cycle (bottom). One thermal cycle comprises a heating half-cycle and a cooling half-cycle. The robot displaces 4.4% body length at the end of the cycle. . . . .	57
<b>Figure 4-2</b>	Plotted on the shape space, phase-lag-induced shape trajectories exhibit the hysteresis often observed in low-bandwidth control systems. The greater the difference in the reaching rate $c_i$ , the larger the phase lag is. This behavior emerges by waiting for a sufficient amount of time for the two joints to both reach steady-state equilibrium and then exciting both joints by a step change in input. . . . .	66
<b>Figure 4-3</b>	Example temperature cycles (our low-bandwidth control input), normalized for period and temperature, where $t_{\text{ramp}} = \eta_{\text{ramp}} \cdot t_{\text{cycle}}/2$ is the time used to ramp input to the goal temperature. Since the ramp time should always be smaller than half of the period, the use of the ramp time ratio avoids additional constraints other than a bounding box. When $\eta_{\text{ramp}}$ is small, the temperature profile is more square-wave-like, and the resulting shape trajectory is more likely to have a larger phase lag. Consequently, for longer ramp times, the shape trajectory is more likely to have minimal phase lag because both joints are in equilibrium for the duration of temperature change. . . . .	67
<b>Figure 4-4</b>	A representative finite element analysis showing unidirectional motion towards the larger bilayer of the thermo-responsive hydrogel crawler subjected to 3 hours of ramped cooling and 3 hours of ramped heating cycle. . . . .	69

**Figure 4-5** Sampled shape data points around a limit cycle. Note that we cannot directly sample the full shape and shape velocity space due to the presence of the actuator dynamics. Fortunately, adding perturbation to the input parameters does generate shape and shape velocity variations around the limit cycle, which is essential for building models of behaviors. . . . . 72

**Figure 4-6** Model prediction metrics for the Purcell swimmer (top left) and hydrogel crawler (top right) model. We plot sample model prediction errors in phase coordinates from each of the model refining iterations. Also, the average displacement, in terms of percentage body length per actuation cycle, each iteration is shown below the plots. The metrics are calculated on the testing set which is unseen by the model both on shape velocity predictions and body velocity predictions. Test set trajectories (grey) along with the mean (solid) and standard deviation (dashed) of the data-driven model prediction error are plotted for each iteration. The body velocity prediction error is inherently propagated from the shape velocity prediction error, and thus the body velocity prediction errors are unsurprisingly greater than the shape velocity prediction error. While the model prediction errors start to drop as we shrink the sampling range due to overfitting, the average behavioral objective improves each iteration. . . . . 73

# Chapter 1

## Introduction

Robotic locomotion is a dynamic and rapidly evolving field that has seen remarkable growth and development in recent years, driven by its vast potential for application in diverse domains. Whether it be in the context of manufacturing [1], healthcare [2], search and rescue operations [3], or even the realm of entertainment [4], the field of robotic locomotion aims to enable agents to master mobility. Wheeled locomotion has been particularly successful in human history thanks to its efficiency and ease of control as a means of transportation. This efficiency comes in part from the minimal dissipation wheeled systems experience on structured environments such as paved roads and train rails. The ease of control is a result of the exploitable symmetries in the interaction of circular shapes on flat surfaces. However, in unstructured and complex environments such as those observed in nature (e.g., sand, grass, rocky surfaces, etc.), wheeled systems are often significantly less effective [5]. In contrast to wheeled locomotion, nature offers a diversity of versatile and robust solutions, such as legged and limbless locomotion, enabling animals to traverse a wide range of environments. This dissertation focuses on topics within the broader field of robotic locomotion that draw inspiration from nature, and seek to provide new methods for data-driven modeling and control of such systems.

## 1.1 Bio-inspired Robotic Locomotion

Robotics researchers have the opportunity to leverage the insights and knowledge acquired from biology in their field. Studying the flight of birds, the swimming patterns of fish, and the walking mechanisms of humans and other species offers valuable lessons. Understanding and abstracting key principles from these biological systems can guide the development of robot designs including legged locomotion (one-legged hopping [6], bipedal [7], quadrupedal [8], hexapod [9]), limbless locomotion [10], swimming [11, 12], and flying [13].

Some robotic designs focus more on the mechanics or functionalities of their inspirations from nature, such as pneumatic-driven crawlers [14] and single-leg hopping robots [6]. Other biomimetic robotic systems aim to resemble biological systems in both the appearance and intrinsic dynamics of movements. Examples include a fish-inspired swimmer [15], a snake-like robot [16], cockroach-inspired RHex family hexapod robots [9, 17], the robo-cheetah [18], and many more. This dissertation primarily discusses locomotion on land, where legged and limbless locomotion schemes are most observed in terrestrial animals including humans. Today, legged and limbless locomotion schemes are of special interest to researchers due to their versatility and robustness in unstructured environments.

Legged robots may have one, two, four, six or more legs depending on their applications. Typically, systems with four or fewer legs have more dedicated and complex legs with control algorithms because they are subjected to stability problems. For example, in bipedal walking, one obvious controlling paradigm is to maintain static stability, which requires the center of mass to be within the support polygon at all times. While it is much easier to implement static walking, it is often slow and energy inefficient [19]. Later, generalizations of the ‘center of mass’ criterion such as zero-moment point (ZMP) [20] and foot rotation indicator (FRI) [21] gained popularity in

generating stable walking movements for robotic systems. Researchers drew inspiration and proposed a reduced-order model called ‘inverted pendulum’, which gave rise to completely un-actuated passive dynamic walkers; this model then was generalized into the ‘spring-loaded inverted pendulum’ (SLIP) model that describes running [22–24]. The SLIP model explains the capability of maintaining stable periodic orbits (or limit cycles) [25] with underactuated or even completely passive legs. From this point, it is believed that underactuation could be leveraged to achieve stable locomotion. Recent efforts in this area include the development of control-enforced virtual holonomic constraints and the use of hybrid zero dynamics [26, 27] to achieve stable walking.

Higher degree-of-freedom locomotion presents new opportunities as well as challenges. Quadrupedal walking is also an important research topic, with many robots developed to mimic the walking patterns of animals such as dogs [8], horses [28], and cheetahs [18]. When it comes to hexapod and multi-legged robots, dynamic stability is of less concern because they typically can maintain static stability with three or more legs on the ground. Note that many of these systems are also dynamically stable for times when they do violate static stability [29]. Hexapod robots with high degrees of freedom legs are capable of precisely planning and controlling each foothold, making it easy to achieve stable locomotion [30, 31]. However, the complexity of the control algorithms and the high number of actuators make these systems challenging to design and implement. In contrast, hexapod robots with fewer degrees of freedom legs are easier to design and implement while having on average more potential power output per leg. For example, the RHex series robots [9, 17] have rotary ‘c’ shaped legs which are only one degree of freedom, while they are still capable of performing various movements, including walking, running, turning, and jumping.

Snakes in nature have incredibly high degrees of freedom which in turn provide plenty of control redundancies to their movements. Researchers observed basic gaits that can be categorized into four types: rectilinear, lateral undulation, sidewinding

movements, and concertina. Rectilinear movement involves alternating movements through the snake’s ribs and muscles to advance the body, making it suitable for narrow spaces. Lateral undulation, characterized by lateral wave propagation, propels the body forward and is ideal for flat terrains. The sidewinding movement, with its spiral characteristic, allows for lateral or oblique movement. This is best exemplified by the gait of desert-dwelling rattlesnakes, enhancing adaptability to the terrain. The concertina movement, seemingly similar to the lateral undulation mode, involves self-tracing to crawl forward and is typically used by snakes when climbing trees. [32] Researchers often use proposed shape curves to control the motion curve of snake-like robots, aiming to approximate the ideal shape curve. For instance, the serpenoid curve proposed by Hirose et al. is used to mimic the lateral undulation mode [33], while Burdick et al. established a three-dimensional motion curve in a piecewise manner to analyze the sidewinding movement [34]. Drawing from these insights and observations, researchers have developed various snake-like robots, such as the 2D passive wheel snake robot [35], ACM R5 snake robot that has shielded joints [36], and a modular snake robot presented in [37]. This motivates methods to procedurally engineer behaviors without needing to rigorously study individual strategies.

Across degrees of complexity, the development of bio-inspired robotics and the study of biological locomotion are heavily intertwined and mutually reinforcing. Efforts are also made the other way around, using robots or physical models as experimental tools to verify current understanding or propose new hypotheses in science. Terms referring to this category include ‘robotics-inspired biology’ [38] and ‘robophysics’ [39]. In this evolving paradigm, researchers are challenging the traditional approach, that exclusively relies on observations, by using robotic and mechanical systems to gain fresh insights into nature [40]. Generally, this approach involves designing physical models and robotic systems, studying their performance and control, and subsequently directing biologists toward new animal experiments and the exploration

of undiscovered biological phenomena.

Within this framework, mechanical and robotic systems become valuable tools for understanding how sensory and feedback systems contribute to the control of intricate biological movements. Additionally, they shed light on how various configurations of actuators power movement and how skeletal elements may impact locomotion. Insights derived from the study of these mechanical systems then guide the exploration of similar phenomena in animals, proposing new hypotheses about animal function that might not have arisen solely from traditional biological studies. This reciprocal relationship between robotics and biology broadens the avenues for scientific inquiry and discovery [41].

Despite the significant progress in this field, there are still many challenges to overcome before robotics can achieve the level of effectiveness seen in biological systems. One of the many challenges is the need for accurate and efficient modeling of these systems, a central aim of this dissertation research.

## **1.2 Modeling Locomoting Systems**

Models play a pivotal role by providing a structured and systematic approach to understanding, predicting, and controlling the behavior of locomoting systems. The underlying physical principles, dynamics, and interactions that govern the movements and operations of robots are typically represented by a small number of simplified, ordinary differential equations. These models serve as essential tools for designers and engineers, allowing them to simulate and analyze the performance of robotic systems in a virtual environment before real-world implementation. By offering insights into system behavior, modeling not only accelerates the design and optimization of robots but also supports their integration into diverse applications, from manufacturing and healthcare to exploration and entertainment. We usually categorize approaches to

modeling locomoting systems onto a spectrum of *bottom-up* and *top-down* approaches, where the former relies on induction from basic physical principles and the latter relies on empirical data.

### 1.2.1 Bottom-up: Induction from First Principles

Towards acquiring such useful models, the more traditional approach has been to first understand the underlying physics and mechanisms, make well-defined and explicit assumptions, and then develop a model in a deterministic manner. This approach is often referred to as a *bottom-up* approach, as it starts with the most basic components and gradually assembles them into a more comprehensive model. By studying the fundamental building blocks and their interactions, researchers can then scale up their models to gain insights into the behavior of larger and more complex systems. Bottom-up modeling is valued for its ability to provide a detailed understanding of a system's inner workings and for its applicability in fields where precision and accuracy are crucial. Such analytical models are beneficial when you have a solid understanding of the system and can make reasonable assumptions about its behavior. On the other hand, when there is a lack of understanding of the underlying mechanisms, or when the system is too complex to be effectively represented by analytical models, constructing *bottom-up* models in such cases can be particularly challenging. For locomoting systems, a mathematical framework *geometric mechanics* has been proven useful in providing structural insights, which plays such an essential role in this dissertation; we briefly review it in the next section.

#### 1.2.1.1 Geometric Mechanics

Geometric mechanics offers powerful tools for studying locomotion [42, 43] that provide deep links between physical motion and fundamental mathematical structures. These methods exploit mathematical symmetry in systems' equations of motion and highlight



relationships between changes in the internal shape<sup>1</sup> of a system and its position in an externally fixed reference frame.

The use of symmetry takes the form of a fundamental assumption that when expressed in body coordinates, the system’s constraints remain invariant concerning its position, including orientation. Intuitively, if a set of changes in a system’s shape is able to move the body frame in certain ways, it will always do so regardless of the system’s orientation or location, as the system’s capabilities to locomote is totally invariant to its position. This symmetry leads to a convenient separation between a system’s shape space and position space, where the combination of both spaces comprises the system’s full configuration space. This separation is formally represented by a mathematical structure known as a fiber bundle, where the shape and position of the system are respectively described as elements of the base and fiber spaces. Connections on this fiber bundle encode combinations of shape and position velocities that conform to constraints on the system, such as the conservation of momentum or interactions with a surrounding fluid.

The early form of the current geometric locomotion framework can be traced back to [45], where it discussed two representations of local nonholonomic constraints arising from angular momentum conservation and body force conservation. Drawing on the mechanical symmetry work in [46], researchers including Kelly, Ostrowski, Burdick, Bloch, and Lewis developed the reconstruction equation [42, 47–49], which generalized the earlier results to apply to systems with first-order nonholonomic constraints, such as non-slip passive wheels. This generalization includes a momentum term reflecting the system’s coasting behavior, facilitating the study of systems falling between completely constrained low Reynolds number systems and those unconstrained except by momentum conservation.

The reconstruction equation found diverse applications in locomotion contexts,

---

<sup>1</sup>The term ‘shape’ refers to a system’s internal configuration [44]. e.g., all of its joint angles.

where it was used with Lie bracket theory to generate sinusoidal gaits for snake-like systems [50]. The reconstruction equation was also used to decouple kinematic and dynamic elements for designing kinematic gaits [51]. The connection vector field was developed to aid in visualizing the reconstruction equation [52]. Recent attention in geometric locomotion has returned to swimming systems, including the formulation of the reconstruction equation for various Reynolds numbers [53–55].

It is typically impossible to integrate the reconstruction equation over arbitrary shape trajectories in closed form. Challenges arise in the context of optimization and planning, where it is difficult to find shape changes producing desired translations. In special cases where the integration is done over closed loops in shape space, Stokes’ theorem can be applied to calculate net translation resulting from gaits [42] by converting the line integral into an area integral. Following this principle, Shamma *et al.* [56] developed *height functions* on the shape space to describe rotation displacements of various gaits for their three-link robot. This development allowed visual design of gaits that result in specified rotations. The above results were limited predominantly to rotations because the noncommutativity of the Lie bracket makes it infeasible to calculate net translations for arbitrary gaits; unless the system has zero turning at all parts of the gait. By carefully selecting the body coordinates, Hatton & Choset [57] showed that the noncommutativity effect on net translations can be greatly reduced. This choice of coordinates can be systematically made so that the body frame has minimal rotational perturbation, making the area integral from the *height functions* much more accurate for translational movements. This procedure was then formalized in differential geometry terms [58].

Some recent efforts focus on applying the geometric paradigm to analyze real-world systems approximated by a local connection. Hatton & Choset [59] demonstrated the utility of geometric tools for providing insights into analytically intractable systems, whose output aligns with the geometric template. Dai *et al.* [60] applied the same

geometric framework to a high degree of freedom snake robot parametrized with two serpenoid mode functions.

On the theoretical side, Ramasamy & Hatton [61] extended the framework to higher dimensional shape spaces where the area integral in [42] is generalized into hyper-surface integral. Choi *et al.* [62] worked on generating gait libraries for planning situations where intermediate gaits are preferred (e.g., smaller steps, making gentle turns). They showed that by carefully selecting the constraint via the construction of a Lagrange multiplier, a continuous gait library that satisfies the constraint can be found through constrained optimization.

### 1.2.2 Top-down: Data-driven Methods

Often, the internal workings or mechanisms of a system are not explicitly known or considered, yet acquiring models for such systems is still important for purposes like control and planning. System identification methods under these scenarios are referred to as *Top-down* approaches. They focus solely on input-output relationships, treating the system as an opaque entity with inputs and outputs, much like a ‘black box.’ Black-box models are often based on empirical data and can rely on machine learning techniques, such as neural networks [63], decision trees [64], and reinforcement learning [65]. They are particularly useful when dealing with systems that are too complex or poorly understood to be effectively represented by analytical models. Black-box modeling has gained prominence in fields like artificial intelligence, where the focus is on predictive accuracy and performance, rather than a detailed understanding of the system’s inner workings. These models make predictions without providing insights into the underlying processes that generate those results. While black-box models can be highly effective at capturing complex relationships in data and making accurate predictions, they lack transparency and interpretability. This lack of transparency can be problematic in applications where the underlying mechanisms are subjected to

changes or modifications, making it challenging to generalize a trained model to new scenarios or slightly different systems [66].

Often, these black-box models consist of large amounts of parameters; and thus require equivalently large amounts of empirical data to capitalize on the complexities of the model. Modern robotic systems involve more and more degrees of freedom for better performance, more redundancies, new mechanisms, and more complex dynamics. The *curse of dimensionality* arises in the context of increasingly complex systems, where higher dimensional action spaces (shape spaces, in our terms) become exponentially more challenging to explore. A common solution is to parametrize the action space where the effective degrees of freedom are reduced to a manageable number. In addition to the *curse of dimensionality*, many systems (and/or their interactions with environments) are subjected to changes and uncertainties due to reasons like wear, manufacturing imperfections, dynamic environments, etc. Although there have been specific techniques developed to mitigate these issues like domain randomization [67], they are often highly specific to the system at hand and require a large amount of data to train. The above complications undoubtedly make pure black-box modeling challenging in terms of data efficiency and generalizability.

### 1.2.3 On the Spectrum of White-box and Black-box

*Bottom-up* (i.e., white-box) and *Top-down* (i.e., black-box) approaches are two ends of the spectrum of modeling techniques, where the former relies solely on induction from basic physical principles, and the latter relies only on empirical data. While both approaches have their advantages and disadvantages, a combination of both approaches, often referred to as grey-box modeling, can be particularly useful. These approaches contain components from both ends of the spectrum. For example, sometimes there are fundamental assumptions that can be made to provide a known structure (say a spring-mass system) and empirical data can be used to fit unknown model parameters

(a spring constant in the case of a spring-mass system). These integrated approaches seek to leverage the strengths of both ends of the spectrum, combining the transparency and interpretability of the bottom-up approach with the flexibility and predictive power of black-box modeling. In this approach, the fundamental structure is modeled explicitly, and data-driven techniques are utilized on top, completing the model. Grey-box modeling is particularly useful when the system’s behavior is not entirely understood, and there is a lack of data to build a purely data-driven model. In such cases, the underlying mechanisms can be modeled explicitly, while the unknown or uncertain aspects can be modeled using data-driven techniques. This approach can be particularly useful in the context of robotic locomotion, where the underlying mechanisms are often complex and poorly understood, and the data available is often limited.

As discussed in 1.2.1.1, geometric mechanics provides a natural white-box component for modeling locomoting systems, where specifically we can make the basic assumption that system behavior is invariant to its location. Empirical data can then be used on top to capture the complex system dynamics. This dissertation specifically builds on this integrated approach. This approach is general enough to apply to a broad class of locomoting systems, yet specific enough to efficiently and accurately capture complex physical interactions from data.

### **1.2.3.1 Data-driven Geometric Methods**

Bittner et al. [68] introduced a data-driven approach aimed at geometric modeling and optimizing primarily kinematic systems. This innovative methodology offers a means of performing *in-situ* system identification for various classes of platforms by locally modeling their behaviors through the observation of approximately repeated behavior cycles. By leveraging the variance in the repetition, researchers obtain regularly sampled time series data that can be effectively fitted to a generalized oscillator [69].

This fitting process assigns an asymptotic phase to all observed samples, facilitating the linearization of connections within phase windows concerning an average behavior or limit cycle, which is determined using the asymptotic phase. Across these phase windows, model coefficients contribute to a Fourier series fit, providing a linearized model that can be sampled as a function of phase.

To facilitate the modeling process, the data samples are grouped into overlapping phase bins, and the choice of bin size is a heuristic decision that can significantly impact the model quality. Larger phase bins offer a broader perspective on the dynamics throughout the entire cycle, reducing the risk of overfitting the data globally. However, they can also result in reduced accuracy within the specific phase bin of interest.

Initially, this approach had limitations in its application to systems with momentum. Subsequent research efforts were dedicated to capturing and modeling systems with dissipating momentum while retaining a first-order prediction structure [70]. Further extensions adapted these models to primarily kinematic systems that lacked high bandwidth control over their shape dynamics. The initial extensions focused on systems with high bandwidth control of a subset of the entire shape space [71], eventually progressing to address systems with low bandwidth control of shape variables [72]. In this dissertation, we explore an extension to the original work [68], enabling real-time updates for *in-situ* modeling and adaptation in the field. Simultaneously, our research contributes to the development of an overarching architecture that can accommodate these extensions into adjacent geometric domains, broadening the applicability and potential impact of this innovative approach.

Zhao *et al.* [73] demonstrated the utility of this framework for modeling a seemingly complex locomotion behavior (multi-legged walking with slipping) which they showed can be modeled with a relatively simple geometric model just like other behaviors such as slithering and swimming. This shows the possibility of efficient modeling of many multi-legged robot designs with this framework [74], further extending the

range of applicable systems. In the vision of practically deploying such robots in the world, considering more uncertainties and unexpected changes, we introduce an adaptive system identification scheme, in Chapter 2, that can be deployed to enhance the real-time capabilities of the modeling framework.

Recently, the framework was extended to address ‘Shape-Underactuated Dissipative Systems’ (SUDS), where the controlled degrees of freedom are restricted to a subset of the full shape (i.e., existing passive joints that are not controllable) [71]. This opens the doors to applying the geometric mechanics framework to soft robots, particularly for robots incorporated high-damping series elastic actuators.

An assumption in [71] is that the controlled degrees of freedom are accessed with high-bandwidth control, meaning we would have full access to the controlled joint velocities. However, such high-bandwidth control is a luxury in cases like stimuli-responsive soft actuators where the actuation dynamics restrict full access to the controlled shape variables. Note that although the actuators generally have inherent low-pass dynamics, it does not exclude the system from being principally kinematic as long as the shape is an intermediate variable that is in the causal chain between actuator dynamics and the system’s movement. Inspired by [75], in Chapter 4, we further extend the modeling framework to systems with only low-bandwidth control over their shape.

### **1.3 Steering and Control of Nonholonomic Systems**

Controlling locomoting systems, especially in the context of robotics, presents practical challenges. Complex systems often require intuitive command inputs that can transform into multi-joint commands, ultimately mapping to the desired outputs. This chapter will delve into the specific needs for controlling locomoting systems, considering the intricate interplay between inputs, control algorithms, and the dynamic

nature of locomotion. As we progress through this thesis, we will explore how these control aspects influence the design and implementation of robotic locomotion systems, providing insights into the complexities involved in achieving robust and adaptive locomotion in robotic platforms.

Unlike bipedal and most quadrupedal vertebrates, insects have legs sprawled outward in fore-aft and lateral directions, which offers them additional stability and exceptional maneuverability in the horizontal plane, making them able to execute very tight turns at high speeds [76]. Researchers [77–80] have made mechanistic models for cockroaches’ various tracking behaviors and found the ability to steer the whole body is essential to achieve efficient tracking. Even with similar sprawled structures, achieving such maneuverability in hexapod robots is challenging. Getting a better understanding of the steering behaviors is an essential part of improving planar maneuverability.

Unlike wheeled systems, many biomimetic systems and animals do not have a trivial steering scheme in terms of differentially driving the joints, especially in artificial systems. Oftentimes, scientists get inspiration from particular biological systems for basic movements, yet robots utilizing these basic policies are far from being as agile as their biological counterparts [81]. There have been studies done in biology to explore which mechanisms or control strategies were used for maneuverability in animals [18, 76]. On the robotics side, there are also efforts made to improve maneuverability and steerability for artificial systems [81]. However, these methods are usually made ad hoc and specific to the system. In Chapter 3, we propose a systematic way of steering complex robotic systems based on the geometric mechanics framework and potentially only need a partially estimated locomotion model of the system.



## 1.4 Contributions

This dissertation aims to address three main practical challenges in various applications of data-driven geometric methods.

Chapter 2 draws inspiration from data-driven modeling, geometric mechanics, and adaptive control techniques. We employ an adaptive system identification framework and demonstrate its efficacy in enhancing the performance of principally kinematic locomotors (i.e., the term ‘principally’ expands the range to systems that are highly influenced by dissipative forces). We showcase the capability of the adaptive model to efficiently accommodate varying terrains and iteratively modified behaviors within a behavior optimization framework. This provides both the ability to improve fundamental behaviors and perform motion tracking to precision. We also demonstrate, via a Purcell swimmer, that our method may be successfully deployed for in-situ behavior refinement, injury recovery, and terrain adaptation, particularly in domains where simulations provide poor guides for the real world.

In Chapter 3, we propose both local and global gait modulation algorithms to modify a nominal gait to provide a single-parameter steering control. Using a simplified swimmer, we numerically compare the two approaches and show that for modest turns, the local approach, while suboptimal, nevertheless proves effective for steering control. A potential advantage of the local approach is that it can be readily applied to soft robots or other systems where local approximations to the constraint curvature can be garnered from data, but for which obtaining an exact global model is infeasible.

The method we present in Chapter 4 constructs a series of connected models comprising actuator and locomotor dynamics based on data points from stochastically perturbed, repeated behaviors. By deriving these connected models from general formulations of dissipative Lagrangian systems with symmetry, we offer a method that can be applied broadly to robots with first-order, low-pass actuator dynamics, including

swelling-driven actuators used in hydrogel crawlers. These models accurately capture the dynamics of the system shape and body movements of a simplified swimming robot model. We further apply our approach to a stimulus-responsive hydrogel simulator that captures the complexity of chemo-mechanical interactions that drive shape changes in biomedically relevant micromachines. Finally, we propose an approach of numerically optimizing control signals by iteratively refining models, which is applied to optimize the input waveform for the hydrogel crawler. This transfer to realistic environments provides promise for applications in locomotor design and biomedical engineering.

## 1.5 Dissemination

The material in Chapter 3 has in full been published in:

- Deng, S., Hatton, R. L. & Cowan, N. J. *Enhancing Maneuverability via Gait Design* in *Proc IEEE Int Conf Robot Autom* (2022), 5799–5805

The material in Chapter 2 and 4 have in large part been submitted for publications and their preprints are available on arXiv:

- Deng, S., Cowan, N. J. & Bittner, B. A. Adaptive Gait Modeling and Optimization for Principally Kinematic Systems. *arXiv preprint arXiv:2310.02141* (2023)
- Deng, S., Liu, J., Datta, B., Pantula, A., Gracias, D. H., Nguyen, T. D., Bittner, B. A. & Cowan, N. J. A Data-Driven Approach to Geometric Modeling of Systems with Low-Bandwidth Actuator Dynamics. *arXiv preprint arXiv:2307.01062* (2023)

The following co-authored publication has largely inspired the work presented in Chapter 4 but does not prominently appear in this dissertation:

- Pantula, A., Datta, B., Shi, Y., Wang, M., Liu, J., Deng, S., Cowan, N. J., Nguyen, T. D. & Gracias, D. H. Untethered unidirectionally crawling gels driven by asymmetry in contact forces. *Science Robotics* **7**, eadd2903 (2022)

The following co-authored preprint is also relevant to but does not prominently appear in this dissertation:

- Shi, R., Chen, K.-L., Fern, J., Deng, S., Liu, Y., Scalise, D., Huang, Q., Cowan, N. J., Gracias, D. H. & Schulman, R. Shape-shifting microgel automata controlled by DNA sequence instructions. *bioRxiv*, 2022-09 (2022)

# Chapter 2

## In-situ System Identification for Principally Kinematic Systems

This chapter was submitted for publication and was made available on arXiv entitled *In-situ System Identification for Drag Dominated Systems* authored by Siming Deng, Noah J. Cowan, and Brian A. Bittner. We re-used the article in this chapter with permission from all the authors.

### 2.1 Motivation

The field of robotics has experienced significant advancements over the past few decades, evolving from systems designed for explicit situations into more generalizable and versatile machines. This transformation has been driven by the ever-growing demands of robots that are capable of interacting with dynamic and unstructured environments. In many such environments, learning from *in-situ* experiences is crucial to functioning in the field. These challenges are particularly evident for systems with high internal degrees of freedom and complex system-environment interactions, where subtle changes in environmental characteristics can drive fundamental differences in emergent behavior.

Current paradigms of motion control are not ideally suited to attaining precision in these unstructured settings. Traditional *bottom-up*, model predictive approaches

to robotic control heavily rely on explicit accounting of physical quantities which can be derived from first principles or hand-crafted to the desired precision of the user [85, 86]. Adaptive system identification extends models to react to anticipated changes in the model. The rate and quality of adaptation are affected strongly by the assumptions about observable perturbations of the model [87–89]. In general, these methods are ineffective when either (a) the range of assumed adaptations is so large that overfitting occurs or (b) the assumptions are too specific such that violations of the assumptions are common. More recent approaches to motion control include reinforcement learning and deep learning approaches, both of which require large magnitudes of training data or infrastructurally cumbersome quantities of real-world data [90–92]. Here, the robot is particularly ill-suited to succeed in unanticipated scenarios, since (by definition) they would not be included in the training set. Efforts in domain randomization have yielded reinforcement learning agents more capable of generalization and sim-to-real transfer, but often at the expense of achieving precise mastery of individual domains [67, 93]. Learning frameworks such dynamic mixtures of experts [94, 95] attempt to achieve robustness and precision with distributed expertise, but the computation for these experts to both decouple from each others’ domains and refine individual expertise has not yet been designed for real-time adaptation in complex fielded environments. Adaptive control, reinforcement learning, and deep learning architectures are not designed to rapidly and precisely handle unanticipated changes to the dynamics.

This limits our ability to advance robots to new application spaces that can handle complexities such as customized body parts, imprecise fabrications, and poorly characterized environments. Such considerations are crucial for soft robotics [96, 97], nano-robotics [98], medical robots [99], and bio-hybrid robots [100], as each is subject to these considerations. In this work, we use an adaptive system identification framework to leverage the *top-down* modeling insights of geometric mechanics [68]. This permits

a framework for adaptation that captures the general space of possible dynamics changes within a class of physics, but it is representationally compact enough to be fit rapidly to data *in situ*. The result enables adapting motion models to distinctly new regimes without knowing *a priori* the nature of that change (e.g. a substrate change, morphology change, or weight distribution change). The assumptions made are that the dissipative nature of the physics persists and that the recent history of the behavior is evolving in a homogeneous environment. A key contribution of our work is that empirical modeling is moved online as the system operates and captures a general class of potentially time-varying factors such as environments and payload changes, affording effective motion control despite unanticipated changes to the robot morphology, weight distribution, or substrate conditions.

We restrict our attention in this work to systems that are principally kinematic [42], a property developed in the geometric mechanics community that describes the locomotion of systems governed by dissipative forces or conservation of zero momentum. These systems enjoy an exploitable reduction from second-order mechanical dynamics to ones that possess a first-order structure, namely a linear mapping from internal configuration velocity to body velocity. Recent work has shown that this reduction provides a reduced computational structure for fitting robot models that are broadly descriptive of many classes of animal and robot behaviors [73, 74]. These works have established that many robot and animal behaviors are describable by the space of dissipative models, which span a dense and broad variety of animal morphologies, weight distributions, and environmental substrates. Bittner et al. [68] introduced a data-driven method for constructing a local model in the neighborhood of an observed gait. This method relies on data points collected from stochastically perturbed and iteratively repeated behaviors. The ability to construct local models makes it possible to sample candidate gaits offline, which provides an opportunity for sample efficient gait optimization as well as gait modulations for continuous gait control [82].

Here, we capitalize on the use of adaptive filters in constructing the aforementioned local models in real-time. The advantages of this reframing are multifold. In this continuously updating paradigm, the model automatically adapts to changes in physical interactions dynamically, rather than awaiting a batch update as in prior work. This actively updated model enables tracking of the model confidence, which is crucial in determining how much empirical data is required to obtain a model with acceptable prediction quality for behavior refinement and control. We show that the rate and quality of adaptation permit refinements of behaviors to master unanticipated environments with speed and precision. We discuss that as internal complexity is added to the system, we encounter situations where engineering decisions between speed and reliable precision need to be made.

Next, we will provide technical background on the adaptive filters, followed by a description of the filter which takes a recursive formulation. We follow this with several experiments that demonstrate the performance of the adaptive filter, noting its accuracy in capturing the dynamics of the principally kinematic Purcell swimmer through a regime change from low drag ratio to high drag ratio swimming. Finally, we offer points of discussion and our conclusions on the performance and implications of the proposed approach.

## **2.2 Background**

The mathematical framework used to capture the motion model of principally kinematic systems is covered here. We also discuss prior efforts to model these systems from data and their current limitations.

### **2.2.1 Principally Kinematic Systems**

The field of geometric mechanics [101] [48] [47] has pursued the applications of symmetries to mechanics. In many cases, these applications result in reduced computational

structures for system representation and analysis. One such insight is the reduced Lagrangian or reconstruction equation, which writes the mechanics of the systems using three distinct equations describing the *generalized momentum*, body velocity, and shape dynamics of the system. Notably, these equations describe the interactions of general second-order systems in ways that one can infer in fundamentally different ways than classical rigid body dynamics formulations. This representation leans on a key symmetry that the environment remains homogeneous, such that, for example, the body velocity equation remains a constant function of shape, shape velocity, and momentum, no matter where you place the object in the environment.

In this work, we will pursue an adaptation scheme that applies to a subclass of these systems. Here, we focus on principally kinematic systems. Commonly, these can be observed as systems where dissipative forces dominates the dynamics [102]. Recent work has suggested this class of dissipation model, despite its simplicity, is highly descriptive of a broad class of robots and animals [73, 74]. Principally kinematic systems are also observed where conservation of zero momentum drives the physics [43]. These scenarios are less commonly observed but applicable to orbital manipulation. In principally kinematic systems, the equations of motion can be reduced, resulting in a shape-defined linear mapping from shape velocity to body velocity. The resulting model is written as (4.1)

$$\xi = -\mathbf{A}(r)\dot{r}, \tag{2.1}$$

where  $\xi$  is the body velocity,  $r$  is the internal configuration, and  $\mathbf{A}(r)$  is the linear mapping termed a *connection*. In situations concerning momentum conservation, we use the phrasing *mechanical connection*, whereas in dissipative regimes the term *viscous connection* is applied.



## 2.2.2 Data-driven methods for geometric systems

Bittner et al. [68] developed a data-driven approach to geometric modeling and optimization of these principally kinematic systems. The approaches provide *in-situ* system identification of various classes of platforms by locally modeling behaviors by observing approximately repeated cycles of the behavior. The variance in the repetition provides regularly sampled time series data that can be fit to a generalized oscillator such that all observed samples are assigned an asymptotic phase. Within phase windows, a linearization of the connection is fit to this data, organized with respect to an average behavior or limit cycle fit using the asymptotic phase. Across these phase windows, model coefficients provide input to a Fourier series fit which provides a linearized model that can be sampled as a function of phase. Perturbations from the average behavior are computed as  $\delta := r - \theta$ , where  $\theta$  provides a shape on the limit cycle as a function of phase. The first-order Taylor approximation of the connection is computed as

$$\mathbf{A}_k(r) \approx \mathbf{A}_k(\theta) + \delta^T \frac{\partial \mathbf{A}_k}{\partial r}, \quad (2.2)$$

where  $k$  denotes rows of the connection matrix.

Samples are grouped into overlapping phase bins. The size of these bins is a heuristic that can strongly affect model quality. Large phase bins give the data set larger insight into dynamics farther along or behind in the cycle. The result of large phase bins is often less overfitting globally, but also less accuracy locally in the phase bin of interest. The regression takes the form

$$\xi_k^{(n)} \sim \mathbf{C}_k + \mathbf{B}_k \delta^{(n)} + \mathbf{A}_k(\theta) \dot{\delta}^{(n)} + \delta^{T(n)} \frac{\partial \mathbf{A}_k}{\partial r} \dot{\delta}^{(n)}, \quad (2.3)$$

where  $k$  provides the output dimension,  $n$  is the sample number. These regressions take place in phase windows centered about a phase-defined shape on the limit cycle,  $\theta_m$ . This model is fitted in phase windows across the cycle and then fitted to a Fourier series, providing a phase-queried linearization of approximately repeated behaviors.

The initial limitations of this approach involved its applicability to systems with momentum. Follow-up work was done to capture and model the behavior of systems with dissipating momentum while maintaining a first-order prediction structure [70]. Additional follow-on work extended these models to principally kinematic systems that do not possess high bandwidth control of their shape dynamics. Initial extensions addressed systems with high bandwidth control of a subset of the entire shape space [71] but eventually worked towards addressing systems with low bandwidth control of shape variables [72]. In this chapter, we explore an extension to the initial work [68] that permits real-time updates for *in-situ* modeling and adaptation in the field. Simultaneously, we work directly towards the overall architecture that can support these extensions into adjacent geometric domains.

## 2.3 Methods

The two core methods of our implementation are presented here. One method involves an adaptive filter for data-driven geometric mechanics that takes a recursive formulation such that it is updatable in real-time for fielded platforms. The second method covers prediction quality metrics, also computable in real-time, which provide guides for when these adaptive models are suitable for trajectory tracking or informing behavior refinements. We also cover a stochastic perturbation scheme for generating approximately repeated behaviors, which is the means by which we generate local behavior data for modeling.

### 2.3.1 Adaptive Filters

Our method employs adaptive filters, which are widely used in system identification. These filters are designed to estimate the parameters of a linear model by recursively updating the weights as new data samples are available. The main usage of adaptive filters in system identification is that they can track time-varying systems and adapt

to changes in the underlying process. Recursive Least Square (RLS) filters are a class of adaptive filters that specifically minimize the mean square error of the predicted filter output and the actual system output (measurements) [103]. Instead of collecting samples in batches and then fitting a model, we designed a cluster of adaptive filters on top of the data-driven modeling framework so that the model can be updated simultaneously as samples are collected. RLS filters fit the role directly in the linear regression problem in (4.3).

In this work, we restrict our attention to approximately cyclic behaviors. For such behaviors, phase is a useful tool through which to organize behavioral data into localized partitions. In phase coordinate, namely  $S^1$ , the nominal gait is written as  $\theta(\cdot)$ ; and phase windows with the same width are set up equally spaced along  $S^1$ . The phase window indexed at  $m$  is defined as  $[\phi^m - \frac{\Delta\phi}{2}, \phi^m + \frac{\Delta\phi}{2}]$ , centered at  $\phi^m$  with a width of  $\Delta\phi$ . The nominal gait at the  $m^{\text{th}}$  window center is denoted  $\theta^m$  with its derivative  $\dot{\theta}^m$ . In each phase window, an RLS filter is constructed for each system output, body velocities  $\xi_k$ . For example, a system operating on SE(2) should have three independent RLS filters in the same phase window, corresponding to its three-dimensional body velocity. Every time a new sample is collected, all filters whose corresponding phase windows cover the sample's phase value are updated as follows:

$$\begin{bmatrix} \vdots \\ \xi_k \\ \xi_k \\ \xi_k \\ \vdots \end{bmatrix} \sim \begin{bmatrix} \vdots \\ 1, \delta^{m-1}, \dot{\delta}^{m-1}, \delta^{m-1} \otimes \dot{\delta}^{m-1} \\ 1, \delta^m, \dot{\delta}^m, \delta^m \otimes \dot{\delta}^m \\ 1, \delta^{m+1}, \dot{\delta}^{m+1}, \delta^{m+1} \otimes \dot{\delta}^{m+1} \\ \vdots \end{bmatrix} \cdot \begin{bmatrix} \vdots \\ \mathbf{w}_k^{m-1} \\ \mathbf{w}_k^m \\ \mathbf{w}_k^{m+1} \\ \vdots \end{bmatrix}^T \quad (2.4)$$

Here,  $\xi_k$  is the desired value of the  $k^{\text{th}}$  body velocity,  $\delta^m := r - \theta^m$ ,  $\dot{\delta}^m := \dot{r} - \dot{\theta}^m$  are the shape and shape velocity perturbation samples defined in the  $m^{\text{th}}$  phase window.  $\mathbf{w}_k^m$  is the weight vector of the  $k^{\text{th}}$  RLS filter in the  $m^{\text{th}}$  phase window, where the dimension is  $1 + 2d + d^2$ ,  $d$  being the dimension of the system shape. The regressor matrix is constructed by stacking the shape and shape velocity perturbation samples

in all phase windows that cover the sample’s phase value. The full model of the system is constructed by combining all the filter weights and smoothed along  $\phi$  by a Fourier series.

An important parameter in the RLS filter algorithm is the forgetting factor  $\lambda_{RLS} \in (0, 1)$ , which gives exponentially less weight to older samples. The forgetting factor, in combination with other parameters (e.g. phase window width and spacing, data sampling rate in terms of phase), determines how much data is effectively in the memory of the model. This number can be tuned for the robustness vs. adaptability trade-off, where larger memory results in a more stable model while being slower to adapt to real changes.

### 2.3.2 Model Metrics

We evaluate the first-order data-driven model by comparing its prediction error to that of the baseline phase model as in prior work [71]:

$$\Gamma = 1 - \frac{\sum_{n=1}^{\mathcal{N}} \|\xi_D^{(n)} - \xi^{(n)}\|}{\sum_{n=1}^{\mathcal{N}} \|\xi_T^{(n)} - \xi^{(n)}\|}. \quad (2.5)$$

$\xi^{(n)}$  is the body velocity from the  $n^{\text{th}}$  collected sample,  $\xi_D^{(n)}$  and  $\xi_T^{(n)}$  denote the data-driven and phase (baseline comparison) model predictions for this velocity. The baseline phase model is effectively the zeroth order approximation, in other words, phase averaged behavior of the seen data.

In an online situation, it is important to keep track of a measure that gives a level of confidence to the model. This moving metric can be used as a criterion for determining whether the model has converged and can be used with confidence. This metric is actively updated as new samples are collected as the following,

$$\begin{aligned} \psi_D^{(n)} &= \lambda_\Gamma \cdot \psi_D^{(n-1)} + \|\xi_D^{(n)} - \xi^{(n)}\| \\ \psi_T^{(n)} &= \lambda_\Gamma \cdot \psi_T^{(n-1)} + \|\xi_T^{(n)} - \xi^{(n)}\| \\ \Gamma &= 1 - \frac{\psi_D^{(n)}}{\psi_T^{(n)}}. \end{aligned} \quad (2.6)$$

Here,  $\lambda_\Gamma \in (0, 1)$  is the forgetting factor, which is set to be nearly 1. This number determines how heavily new samples are weighted towards the measure and how fast the history of the metric fades away. Similar to the forgetting factor  $\lambda_{RLS}$  in the RLS algorithm,  $\lambda_\Gamma$  (together with other parameters) effectively controls how much history the metric  $\Gamma$  is carrying. For convenience, we will later refer a relatively ‘rapid’ forgetting factor to  $\lambda$ ’s further away from 1, and a ‘slow’ forgetting factor to  $\lambda$ ’s closer to 1. In practice, we try to match this metric memory to the effective model memory, in terms of the number of cycles of empirical data, so that it reflects model qualities on the same amount of data. Note that, in doing so, the two forgetting factors are going to be different because each RLS filter updates only when the sample phase falls into its window while the metric is being updated with every sample collected.


### 2.3.3 Online Noise Generation

The data-driven modeling framework needs to have good coverage of both shape and shape velocity around the nominal gait. Hence, we need to generate perturbations in the neighborhood of the nominal gait while commanding the system in real-time.

To ensure the relative smoothness of the generated shape trajectory, we employed a second-order stochastic differential equation (SDE) as a smoothing filter,

$$d\dot{\delta} = -(\alpha\dot{\delta} + \beta\delta)dt + \eta \cdot dW \quad (2.7)$$

where  $\eta$  is the brownian noise amplifier coefficient,  $\alpha$  and  $\beta$  are the attraction coefficients that are kept positive, and  $dW$  is the Wiener process. The noise generated by this SDE is a smoothed first-order perturbation on the system shape velocity, which also ensures that the resulting shape trajectory is pulled back to the nominal gait over time.



**Figure 2-1.** The Purcell swimmer incrementally experiences more stochastically perturbed cycles in a new operating environment. Model metrics of the batch model (green) and adaptive model (red) are evaluated for their predictive performance at that point in the experiment on a hold-out test set of 40 cycles (for 100 pairs of training and testing trials). The model metrics for both methods are shown in boxplots where each box corresponds to 5, 25, 50, 75, 95 percentiles of a model trained with the first  $n \in [5, 10, 15, 20, 25, 30, 35, 40]$  cycles of data experienced.

## 2.4 Accuracy of Adaptive Geometric Modeling

We first seek to demonstrate the speed and precision through which RLS can provide geometric models of locomotion in real time. We demonstrate the method’s performance first on a toy locomotion system, the Purcell swimmer [104], whose local connection can be analytically derived [55, 59]. This analytical system provides a test bed for our method that has clear foundations and has been well explored in the robotics literature. As noted earlier, small soft swimming machines are included within the principally kinematic operating environment for which these methods are designed.

Selecting the right data size for the batch model is important because too large a data size is unnecessary which leads to data inefficiency, while too small a data size is simply not enough to capture the dynamics of the system. From trial and error, batch modeling a nine-link Purcell swimmer typically requires at least 30 cycles of data. However, this number is not easily generalizable for arbitrary systems. One advantage of the adaptive model is that it can keep track of the model’s accuracy in real time so that there is no need to guess a data size as in the batch model. The tracing accuracy (see (2.6)) can help determine whether the model has converged and thus make the adaptive models more efficient in data usage.

We compared our adaptive models to the batch model mentioned in 2.2.2, in terms of model metrics  $\Gamma$  covered in 2.3.2. Again,  $\Gamma$  captures the knowledge of the model beyond 0 (the phase-averaged behavior) up to 1 (full reconstruction of the ground truth). We generated 200 trials of 40 cycles of noisy shape trajectories as mentioned in 2.3.3. We split the data in half, selecting two 40 cycle trials at a time (100 pairs) and evaluating the model computed on the training set against the held out 40 cycle test set. In Fig. 2-1, the model accuracies are plotted with respect to the robot’s experience throughout the trial. This informs us how the prediction quality grows as the robot’s experience grows. The model metrics are shown in boxplots where each box corresponds to 5, 25, 50, 75, 95 percentiles of a model trained with the first  $n \in [5, 10, 15, 20, 25, 30, 35, 40]$  cycles of data collected and tested on the holdout set. When learning from no history, the adaptive model is capable of obtaining an accuracy of  $\Gamma > 0.4$  in 10 cycles, meaning that it can understand forty percent of all perturbations from the limit cycle. This nears the approximate convergent prediction quality of both methods after 40 cycles. The adaptive model metrics are computed continuously, and so are available in between the discretely updated batch model.

As shown in Fig. 2-1, the adaptive model performs relatively well even when there are relatively few samples, whereas the batch model performs best with large datasets.


Indeed, the adaptive method generally affords faster convergence yet a lower overall accuracy than the batch method. With relatively small data sizes, the batch model suffers from near-singular covariance matrices in the local linear regression processes, which leads to poor performance. On the other hand, the adaptive model naturally avoids this problem since the RLS algorithm does not require matrix inversions. Eventually, the batch model converges to a higher accuracy than the adaptive model, which is expected since the batch model has access to all the data while the adaptive model only has a limited memory size. The variation in model quality over the duration of the 40 cycles for the adaptive method speaks to the overfitting produced by the recency bias of the measurements.

Forgetting factors are designed to help weight current (and possibly different) domains more highly, but will also lead the model to potentially overweight data that it has seen more recently. The batch model considers all samples equally, and it is thus less prone to overfitting. Next, we revisit the forgetting factor as a key enabler for the swimmer to regain predictive power on new, unanticipated substrates Fig. 2-2.

## 2.5 Adaptation to Changing Environments


In the derivation of the Purcell swimmer’s connection matrix (which defines its mapping from shape velocity to body velocity), the ratio of lateral and longitudinal drag coefficients [105] directly affects the nature of locomotion. Typically, this ratio is assumed to be larger than one because the swimmer links are assumed to be slender such that drag is stronger in the lateral direction of link motion. When the drag ratio is 1, the system is unable to exploit the asymmetry of drag to create motion and is unable to displace its center of mass. When the drag ratio approaches infinite, each link acts like a non-slip wheel, allowing it to move farther per cycle than when the drag ratio is classically set to 2.



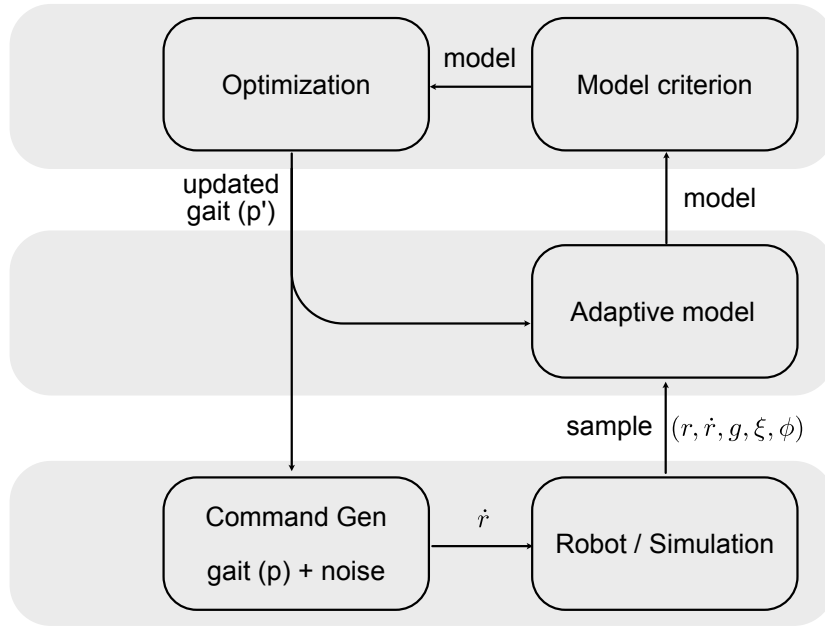


**Figure 2-2.** The predictive quality of the adaptive models (blue,  $\lambda_{RLS} = 0.99$ , and orange, ( $\lambda_{RLS} = 0.7$ )) and batch model (green) are compared throughout environmental changes experienced by the Purcell swimmer. Each model is trained with 40 cycles of data collected in ‘Env 1’. In this test, the system is unknowingly exposed to two other environments, where each environment consists of 40 cycles of data. We changed the drag coefficient ratio experienced on the links for the environment changes, where this ratio is 2.0, 3.0, 4.0, respectively for environments 1, 2, and 3. The batch model (which is never recomputed during the trial) has constant large predictive errors on new substrates. The adaptive model is refined each time a sample is collected and can adapt to the new environments relatively quickly. We note that the forgetful model has a faster adaptation rate but achieves lower, less stable prediction quality throughout the trial.

By altering this drag ratio, we are able to change the physical interactions and test our adaptive method. We trained the batch and adaptive models with data sampled at the classical drag ratio and then adjusted the ratio over a series of interventions. We track the speed and precision with which each model adapts in Figure 2-2. The log prediction error shows that the adaptive model is able to converge to a higher accuracy steady state prediction and does so more quickly than the batch model. The forgetting factor modulates the rate of adaptation, but in this trial, there appears to be a relationship between adaptation rate and overfitting of the model, which is



**Figure 2-3.** The predictive quality of the adaptive models (blue,  $\lambda_{RLS} = 0.99$ , and orange, ( $\lambda_{RLS} = 0.7$ )), batch model (green), and ground truth model before injury (red) are compared throughout an injury (links broke off) experienced by the Purcell swimmer. Each model is trained with (a. 50, b. 100) cycles of data collected in normal conditions. In this test, the system originally consists of (a. 4, b. 9) links, and it undergoes an injury, where (a. one, b. two) links break off. The batch model is not updated within the trial and it has constant significant predictive errors in the new conditions. The adaptive model is refined each time a sample is collected and can adapt to the injury relatively quickly for both low and high dimensional systems. The ground truth model is the simulation model that generates the data used for this demonstration, hence the log error of this model before the injury is negative infinity. It quickly degrades after the injury happens.



**Figure 2-4.** Above we lay out an optimization architecture for behavior refinement. (Top Level) The optimization will attempt to improve behaviors only when it receives a high enough prediction confidence from the adaptive model. (Middle Level) The adaptive model recursively updates from real-world interactions, also producing an estimate of its model confidence in real-time (quantified as  $\Gamma$ ). (Bottom Level) A nominal gait, parameterized as  $p$ , is given to the *Command Gen* module, which generates stochastic perturbations around the nominal gait.

consistent with our discussion of forgetting factors in 2.4.

## 2.6 In-situ Gait Optimization with Adaptive Geometric Modeling

Here we show that sample-sample refinements of the adaptive model allow for accelerated exploitation of policy gradients for rapid behavior optimization. A gait optimization in the experiments considered starts with an initial nominal gait, where the system is commanded to move in the neighborhood of the nominal gait using noise signals generated online (see sec. 2.3.3). In past work, gait optimization for kinematic systems involved collecting a 30-cycle model, building a batch model (identical to that used as a baseline in Section 2.4), and using that model to step along a policy gradient

in the gait parameter space. 30 iterations typically resulted in convergence for the nine-link Purcell swimming totaling 900 total cycles [68]. In this work, we update the adaptive model until we reach an accuracy threshold of  $\Gamma = \frac{1}{2}$  (taken over all prior cycles of the current iteration), and then step along the gradient of the policy to select the next set of gait parameters to evaluate. By this architecture, we can ensure that we use the adaptive models to refine the behavior as soon as they provide an accurate landscape of the local dynamics. Since each iteration of the optimization only locally changes the dynamics, we expect the RLS filters to smoothly recalibrate the model weights to adapt to the shifted domain of the shape space. In particular, we managed to obtain faster convergence by updating the zeroth order model (the first element in each RLS weights  $\mathbf{w}[0]$ ) using previous model predictions on the new gait.

$$\mathbf{w}'_k[0] = [1, \mathbf{0}, \mathbf{0}, \mathbf{0}] \cdot \mathbf{w}'_k = [1, \Delta\theta, \Delta\dot{\theta}, \Delta\theta \otimes \Delta\dot{\theta}] \cdot \mathbf{w}_k \quad (2.8)$$

Here, the subscript  $k$  is the  $k^{\text{th}}$  body velocity,  $\Delta\theta = \theta' - \theta$  is the difference between the new and old nominal gaits, and  $\mathbf{w}'_k[0]$  is the first element of the model weights (which captures the phase averaged behavior of the new nominal gait). The rest of the weights are left unchanged because we assume that the new nominal gait shares some variational relationships with the prior gait due to their proximity in the shape space. This provides a seeded model which we assume will require less refinement for the RLS adaptation. This zeroth-order term update is equivalent to a Taylor polynomial rebase. The full optimization pipeline is shown in Fig. 2-4.

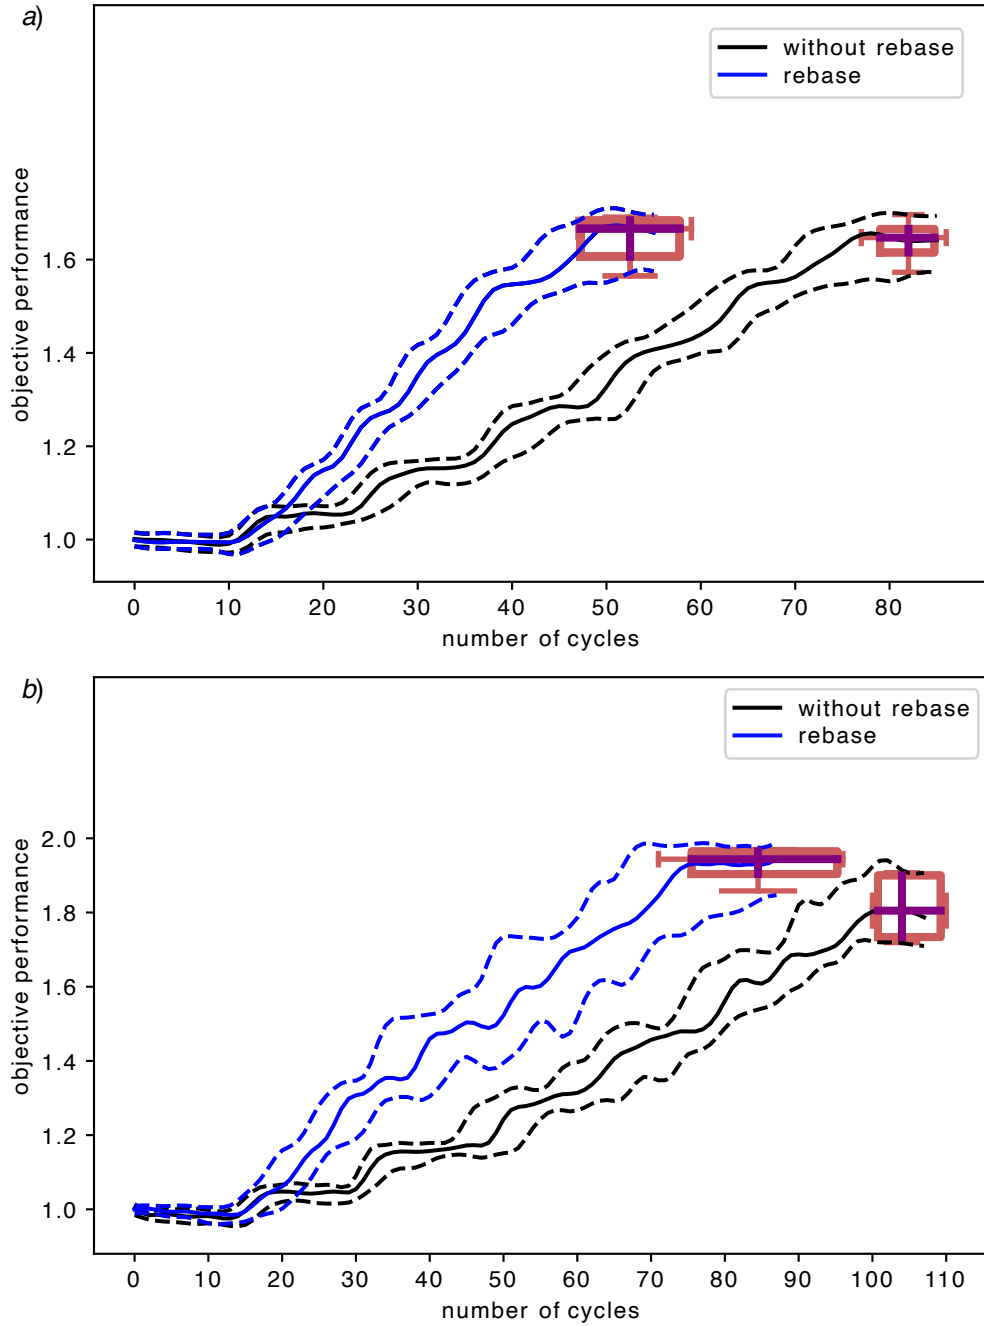
We further demonstrate the effectiveness in data efficiency of rebasing local linear models in 2-5. The adaptive model is effective at capturing unanticipated changes. In the context of changing nominal gaits during the optimization process, the robot should be able to anticipate changes to the model because it is actively changing the gait that it operates around. Thus, rebasing the local model weights can utilize the current model, which comes from previous data points, to obtain an estimate for the new gait. In this experiment, we compare the performance of the adaptive

model with and without linear rebasing. The metric criterion for optimization is set to  $\Gamma = 0.8$ . The adaptive model with rebasing is able to achieve the same level of model accuracy with fewer empirical data, which leads to faster convergence of the optimization process.

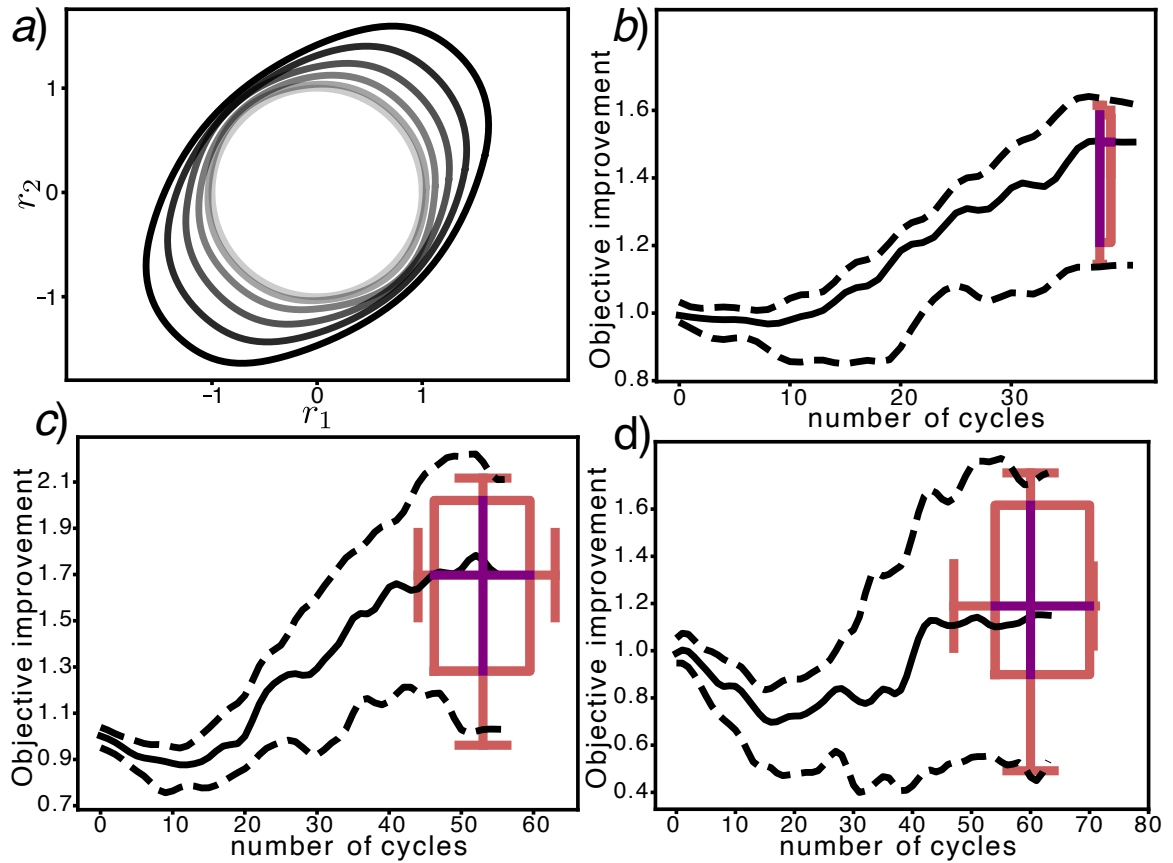
Figure 2-6 covers an application of this architecture to the Purcell swimmer. We repeat the optimization 50 times such that we can observe the range of performance and the repeatability of the process. Each swimmer is seeded with an initial gait that is a first-order Fourier series with equal phase lags between each link. The three-link swimmer is capable of quickly accessing  $\Gamma = \frac{1}{2}$  predictive insight and continuously refines its gait over the 40 cycles, improving its seed behavior by almost 50%. The nine-link swimmer, while less reliable than the three-link swimmer, experiences trials in its top quartile that exhibit over 80% improvement of the gait in just 60 cycles. Our proposed method appears to exhibit a 10-factor improvement in sample efficiency over the work of [68]. Notably, the adaptations are not exclusively successful on the nine-link system, as they were in prior work. For more complex systems, careful consideration of appropriate  $\Gamma$  and  $\lambda_{RLS}$  may improve the reliability of optimization at the expense of sample efficiency.

## 2.7 Discussion and Conclusion

In this chapter, we illustrate the real-time capabilities of our adaptive methods for constructing local models for *in-situ* motion control and behavior refinement. One immediate advantage of integrating adaptive features is the model’s capacity to autonomously adjust in real-time to unanticipated events such as variations in terrain or payload. The prior batch modeling was unreliable at achieving this real-time capability due to the lengthy requirement to run a batch update at every iteration. This real-time enhancement brings fast, precise *in-situ* modeling to the field, which has classically been a challenge for platforms with poor first principles models as well as



**Figure 2-5.** An illustration of rebasing the local linear model at the nominal gait change can further improve data efficiencies for (a) 3-link and (b) 5-link Purcell Swimmer in gait iterations. For no-rebase trials (black), we let the model adapt its weights to reflect nominal gait changes. For rebase trials (red), the model weights are updated using the previous model predictions on the new nominal gait as described in (2.8). The rebase method is able to achieve similar objective function values with fewer empirical data.



**Figure 2-6.** Gait optimization results for Purcell swimmers using the adaptive model optimizer. (a) The three-link Purcell swimmer can refine its gaits from grey to black, optimizing in 40 cycles of experience. 50 trials of the (b) three-link, (c) five-link, and (d) nine-link Purcell swimmer optimization for extremality, with current progress plotted throughout the experience of the trial. Relative performance in displacement per cycle is shown compared to the initial behavior, with [5,50,95] percentiles of performance plotted throughout and [5,25,50,75,95] percentile boxplots provided for the end of the run. The width of the boxplot indicates the number of cycles executed before reaching the final performance value.

systems that encounter unanticipated scenarios (i.e., not accounted for in training sets). Unlike reinforcement learning agents and finely calibrated, explicit physics models, our adaptive scheme is designed to refine models quickly and accurately with no assumptions about morphology, weight distribution, or environmental characteristics (so long as the behavior remains principally kinematic).

We demonstrated this by showing the ability of the adaptive RLS technique to recapture an accurate physical model under significant substrate variation. The adaptation outperformed prior work in terms of speed and data efficiency on the Purcell swimmer as shown in Figures 2-1 and 2-2. These results indicate that the proposed method is an asset for attaining effective motion control in dynamic environments. Also, the adaptive model provides nearly instant reactivity in terms of model updates, which is crucial for real-time applications, while the batch model has a much longer latency in updating the model (mainly due to the matrix inversions in all the local linear regressions).

We noted the relationship between forgetting factors, the rate of adaptation, and the consequences of overfitting as observed in our results. Expectedly, rapid forgetting factors yielded the fastest response to new substrates but the least stable performance. Having a dial to tune can be crucial in settings where either (a) speed is critical to success and risk is tolerable (b) speed is not important and safety is paramount.

We presented an optimization architecture where the system executes stochastically perturbed repetitions of the same gait until it gains requisite confidence in its data-driven model of the local dynamics. When this confidence is met, it leads to high-frequency steps along a computed policy gradient, using this data-driven model to select new gaits that locally improve some performance criteria. This led to a significant performance boost for the state of the art, where the Purcell swimmer optimized its gait to an 80% improvement within 60 cycles in its top quartile of performance. It is likely that a large range of the swimmer behavior is attributed to a high forgetting



factor  $\lambda_{RLS}$  and admissible model prediction threshold  $\Gamma$ , which could interact to give false confidence, resulting in suboptimal steps along the policy gradient. Principled selection of these parameters will likely yield different outcomes, such as a less rapid optimizer that converges with near certainty to a similar behavior. This will be the subject of further research.

The gait optimization result presented here demonstrated the method’s adaptability to gait changes, involving shifts in the sampling region. We note that the rebasing of our model coefficients across discrete changes in the behavior space was crucial to accelerating the model adaptation rate. Notably, we acknowledge that it could also reinforce an overfit model. However, across many trials, this rebase accelerates the model’s capability to explore uncharted regions in the gait space, enabling systematized integration of control, modeling, and optimization processes in real-time with varying physical interactions. In future work, we seek to integrate these models into model predictive control architectures [106–108] to demonstrate refined trajectory tracking that persists through unanticipated scenarios.

# Chapter 3

## Enhancing Maneuverability via Gait Design

This chapter was previously published in IEEE International Conference on Robotics and Automation 2023 (ICRA) as a contributed paper entitled *Enhancing Maneuverability via Gait Design* authored by Siming Deng, Ross L. Hatton and Noah J. Cowan. We re-used the article in this chapter with permission from the publisher and all the authors.

### 3.1 Motivation

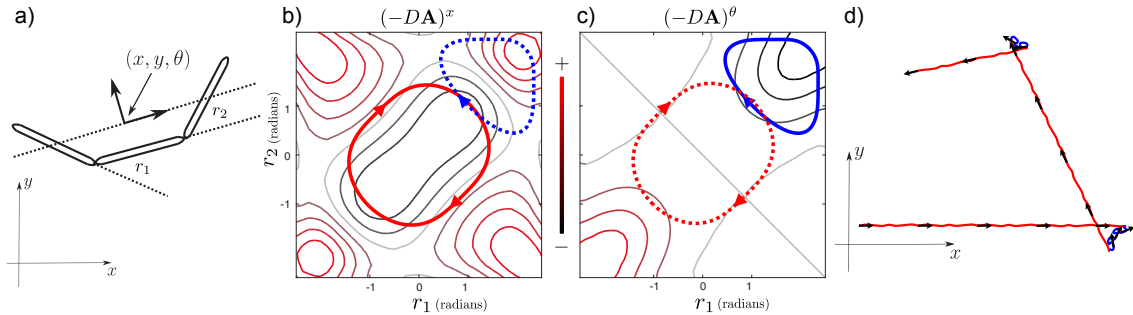
Often, the focus in locomotion research is on the design of a nominal gait—for example, to minimize energy or maximize speed. However, maneuvering and steering are essential aspects of controlling a system. Though switching between a “forward” gait and a “turning” gait would provide a simple solution, continuous steering is crucial to allow a graceful means to respond to sensory feedback, avoid obstacles, or track moving targets. The value of such a graded steering control parameter becomes obvious when one considers driving a car, which would be awkward (at best) if the car only produced a finite set of discrete turning radii from which a driver (or algorithm) must select at each moment in time. By contrast, it would clearly be desirable to be able to exert continuous control over the steering direction. However, most complex

undulatory robots—such as snake robots [33, 109] and legged systems [9], have no such built-in “steering wheel.”

The problem we address in this chapter is how to efficiently modulate an optimized forward gait to produce graded steering control, and effectively nudging the forward gait into a turning gait in a smooth fashion, analogous to a steering wheel in a car. Indeed, cockroaches exhibit graceful steering, as they modulate their gait to “servo” along walls, using tactile feedback from their antennae [110]. Inspired by cockroaches, engineering experts have devised smooth, but *ad hoc* steering control inputs for hexapods such as RHex [9] and Sprawl [80], by effectively inducing graded left–right asymmetries in the gait. The goal of this work is develop a principled way of constructing and exploiting such asymmetries for control.

The present work builds on a massive literature on maneuverability, but we focus narrowly on a few prior works most related to the present paper. Early work [33, 111] addressed how a multi-link snake-like robot that performs undulatory locomotion can make turns by offsetting its body curvature (i.e., steering can be achieved by regularizing the offset level), inducing a lateral asymmetry, a general idea later formalized in a seminal paper on steering control for cyclic, nonholonomic systems [112]. More recent work [113, 114] provided an approach to create gait libraries using a geometric mechanics framework, establishing the framework for maneuverability we build on in this chapter.

Here, we aim to incorporate a previously published geometric optimization algorithm [115] to create general principled means by which to create effective, efficient modulations of a nominal gait for undulatory, drag-dominated robotic systems. If the locomotion model is a local, data-driven model [68]—only available in a neighborhood of the nominal gait—we present a computationally simple, local-gradient-based approach. But if the locomotor model is global (e.g., built from first principles), we can apply an (approximately) globally optimal approach. Then we compare the



**Figure 3-1.** An overview of an example system and its optimal gaits. (a) Schematic of a three-link viscous swimmer—a minimal template [117] for understanding locomotion. The shape of the system, described by its two joint positions, lives in a two-dimensional space. The coordinate of the body frame is chosen according to the *minimal perturbation* principle [57]. (b) The optimal gait (red) in the sense of  $x$ -direction motion efficiency (i.e. the largest displacement in  $x$ -direction per unit power dissipation). The gait is plotted over the system’s *constraint curvature* height function corresponding to  $x$ -direction movement  $(-DA)^x$ , which provides an approximation to the system’s locomotor behavior. (c) Similar to (b), the optimal gait (blue) for turning motion. The two optimal gaits are cross-illustrated in both (b) and (c) with dashed curves. (d) A trivial demonstration of the system’s spatial trajectory from a gait-switching execution. The black arrows over the trajectory denote the system’s  $x$ -axis at the beginning of each cycle, and the color-coded trajectory represents which gait the system is operating on.

performance between the approaches through numerical simulations in a simplified swimmer.

## 3.2 Background

this chapter will use Purcell’s three-link swimmer [116] as an example system to test our gait modulation methods (Fig. 3-1). The three-link model is widely adopted because its simple dynamics can be easily visualized in a two-dimensional shape space.

### 3.2.1 Locomotion Model

When analyzing a system’s locomotor behavior, it is convenient to separate its configuration space  $Q$  into its shape space  $M$  and position space  $G$ , namely  $Q = G \times M$ . A system’s shape (i.e., its joint configuration)  $r \in M$  defines its internal configuration—namely, where each body segment is located with respect to a body-fixed frame.

Meanwhile, its global position  $g \in G$  ( $SE(2)$  for a planar system) locates the system in a spatial frame, specifying the location and orientation of the body frame. The three link swimmer is considered a planar system with a two-DoF shape space, namely  $g \in SE(2)$  and  $r = [r_1, r_2]^T$ . In the sequel, when convenient, we will conflate  $g$  with its (global) coordinate representation, namely  $g = [g_x, g_y, g_\theta]^T$ .

For a drag-dominated system (i.e., a system with negligible inertial effects), the body velocity is related to its shape velocities by the local connection [42, 118]:

$$g^{-1}\dot{g} = -\mathbf{A}(r)\dot{r}, \quad (3.1)$$

where  $g^{-1}\dot{g}$  is the body frame velocity of the system,  $r \in M$  denotes the system's internal shape variable, and  $\mathbf{A}(r)$  is the local connection matrix. The local connection matrix  $\mathbf{A}(r)$  contains all the information of the system's locomotor behavior in the environment; it acts analogously to a Jacobian in that it maps velocities in one coordinate to velocities in another (however, unlike a Jacobian matrix, it is generally non-integrable). The net spacial displacement of performing one gait can be calculated with a path integral over the local connection vector fields:

$$g_\phi = \begin{bmatrix} g_x \\ g_y \\ g_\theta \end{bmatrix} = \oint_\phi -g\mathbf{A}(r) \quad (3.2)$$

where  $g_\phi$  denote the net displacement generated executing one cycle of gait  $\phi$ . Previous works [58] show that we can simplify this line integral by a surface integral as follows:

$$g_\phi \approx \exp\left(\iint_{\phi_a} (-D\mathbf{A})\right) \quad (3.3)$$

where  $(-D\mathbf{A}) = d\mathbf{A} + \sum_{i<j}[A_i, A_j]$  is the total Lie bracket of the local connection. Here,  $d\mathbf{A}$  is the exterior derivative, and  $\sum_{i<j}[A_i, A_j]$  is the sum of Lie brackets between the columns of  $\mathbf{A}$ .

Another important aspect to a drag-dominated system is the effort required when changing shape, which can be modeled as the path-length  $s$  of the system shape

trajectory weighted by a Riemannian metric  $\mathcal{M}$

$$ds^2 = dr^T \mathcal{M} dr, \quad (3.4)$$

that encodes the actual effort required to move the links through the surrounding fluid [116, 119, 120].

### 3.2.2 Gait Optimization

Given a gait parametrization  $p$ , optimal cycles must satisfy the following gradient condition [115]

$$\nabla_p J = \mathbf{0} \quad (3.5)$$

where  $J$  is an objective functional. For example, we consider

$$J_x = \frac{g_x}{s} \quad (3.6)$$

as the objective functional that captures how efficient the gait is in terms of moving forward, regularized by its path-length. Similar objective functionals can be defined in terms of rotation or lateral motion. Solutions to (3.5) for the cost (3.6) can be reached by finding the equilibria of the dynamical system

$$\dot{p} = \nabla_p J_x = \frac{1}{s} \nabla_p g_x - \frac{g_x}{s^2} \nabla_p s + \nabla_p \sigma, \quad (3.7)$$

where  $g$ ,  $s$  denote the spatial displacement and the path-length cost, executing one cycle of the gait  $\phi$ . The other terms in the expression will be described below.

The first term  $\nabla_p g_x$  alone would lead to a gait that has the greatest displacement per cycle, corresponding to the simple objective  $g_x$ . The approximation in (3.3) is a surface integral whose boundary is defined by the gait  $\phi$ ; hence a variation around gait  $\phi$  is intuitively thought of as adding or subtracting weighted regions near the surface boundary. Formally, the general form of the *Leibniz rule*<sup>1</sup> converts this gradient of the functional with respect to variations of its boundary to the gradient of its boundary

weighted by the integrand:

$$\nabla_p \iint_{\phi_a} (-D\mathbf{A})^x = \oint_{\phi} (\nabla_p \phi) \lrcorner (-D\mathbf{A})^x. \quad (3.8)$$

Specifically, the sign  $\lrcorner$  denotes the *interior product*<sup>2</sup> between the boundary gradient  $\nabla_p \phi$  and the weight  $(-D\mathbf{A})^x$ . In the simple setting that we consider here, a system with a two-dimensional shape space, the above interior product reduces to a normal multiplication between the normal component of  $\nabla_p \phi$  and the scalar magnitude of the constraint curvature,

$$\nabla_p \iint_{\phi_a} (-D\mathbf{A})^x = \oint_{\phi} (\nabla_{p\perp} \phi) (-D\mathbf{A})^x. \quad (3.9)$$

Running the optimization with only the first term will end up with gaits located on the zero contour of the constraint curvature.

The second term  $\nabla_p s$  in (3.7) measures how variations of gait affect its execution cost per cycle. This term acts to regularize the optimization, cutting off diminishing returns around low-yield regions on the constraint curvature map. The third term  $\nabla_p \sigma$  in (3.7) is responsible for reorganizing waypoints in the local tangent direction to achieve the most efficient pacing when executing the gait. This term applies changes that keep the gait within the same image-family<sup>3</sup>. For the dissipative mechanical systems that we consider here, this gradient term is orthogonal to the other two terms mentioned above; this implies that the optimization of efficient pacing is independent from the path optimization. In other words, the gradient responsible for pacing,  $\nabla_p \sigma$ , does not affect either displacement per cycle or path-length cost (i.e., it is orthogonal to both  $\nabla_p g_\phi$  and  $\nabla_p s$ ).

For the example system that we consider for this chapter, the optimal gaits, that are generated by previously discussed algorithms, maximizing  $x$  and  $\theta$  motion

---

<sup>1</sup>See [121] for more details

<sup>2</sup>An operation between a vector field and a differential form that results in a -1 degree differential form by "pre-specifying" the vector field as one of the inputs to the differential form. See more in [121].

<sup>3</sup>The image-family of a gait is the set of all gaits that share the same image in the shape space

efficiencies are shown in Fig. 3-1, respectively. The optimal gaits each encircle a rich, sign-definite region, where it gives up the low-yield regions for shorter path lengths, on their respective constraint curvature maps.

### 3.3 Maneuverability

A common scenario in locomotion is steering during forward motion—e.g., a robot executing a translational gait, and modulating this nominal gait to accomplish graded rotational redirection of the heading. Ideally, such steering occurs while largely preserving the overall translational motion. In essence, we are looking for a smooth transformation from the optimum under one objective (e.g. pure forward motion) to the optimum under another objective (e.g. pure rotational motion). Depending on how much information we know about the model, we can generate such transformation of gaits with local or global approaches. In the following subsections, we discuss the local and global methods respectively under the setting of enhancing turning around a forward-moving gait.

#### 3.3.1 Gradient-Based Approaches

The gradients of displacement and cost with respect to the gait parameters provide locally optimal changes with respect to the new objective. Tweaking the gait along the gradient direction will produce the fastest changes in the outcome; in other words, the modulated gait will have significant changes under the new objective with minimal modulation.

Under the previously mentioned setting, to modulate around an optimal forward gait in order to achieve turning, we initiate at the optimal gait in the sense of the most efficient  $x$  displacement as mentioned in 3.2.2:

$$p_x : \nabla_p J_x = \nabla_p \frac{g_x}{s} = \mathbf{0}. \quad (3.10)$$



Modulating this starting gait along the gradient  $\nabla_p J_\theta$  will provide the most rapid changes of  $J_\theta$ , rotational efficiency. Then, the solution of the following equation will guide the gait toward the rotational optimum:

$$\frac{dp}{du} = \nabla_p J_\theta, \quad p(0) = p_x, \quad (3.11)$$

where  $p_x$  denotes the optimal forward gait parameters. The solution  $p(u)$  can be pre-calculated and stored as an infinite-dimensional gait library where  $u$  serves as the modulation variable.

Calculating the above solution will require the model information (i.e. the constraint curvature) at each iteration during the process; visually it is the shape space area swept by the gait library. In situations where we do not have the luxury of knowing the full model, as mentioned in Section 3.1, it is viable to simply use the first-order approximation as the following:

$$p(u) = p_x + u \nabla_p J_\theta|_{p_x}, \quad (3.12)$$

where  $\nabla_p J_\theta|_{p_x}$  is the first gradient at the x-direction optimum. As a comparison, the first order approximation requires less model information as well as less computational power while maintaining similar performance within some range of turning behavior (see Results, Section 3.4).

### 3.3.2 Pareto Front

Now assume we have the full locomotion model, and we would like to push for the optimal transition from one optimum to another. Here we consider the trade-off between  $x$  and  $\theta$  efficiencies in the sense of Pareto optimality, where we cannot improve one objective without hurting the other. A set of Pareto optimal gaits can be generated, forming a Pareto front, using the following objective in our optimization algorithm while iterating  $\alpha$  from 0 to 1:

$$J_{total} = \alpha \frac{J_x}{J_{xmax}} + (1 - \alpha) \frac{J_\theta}{J_{\theta max}}, \quad (3.13)$$

where  $J_x$  and  $J_\theta$  are normalized using the optimal values in each direction alone  $J_{xmax}$   $J_{\theta max}$ .

Although the above method can guarantee global optimality to the objective  $J_{total}$  without being trapped at local optima (i.e., critical points) like the gradient-based methods, it cannot guarantee a smooth transition among gaits located on the Pareto Front. We will continue the comparison between this global method and previous gradient-based methods in Results (Section 3.4).

### 3.3.3 Online Control

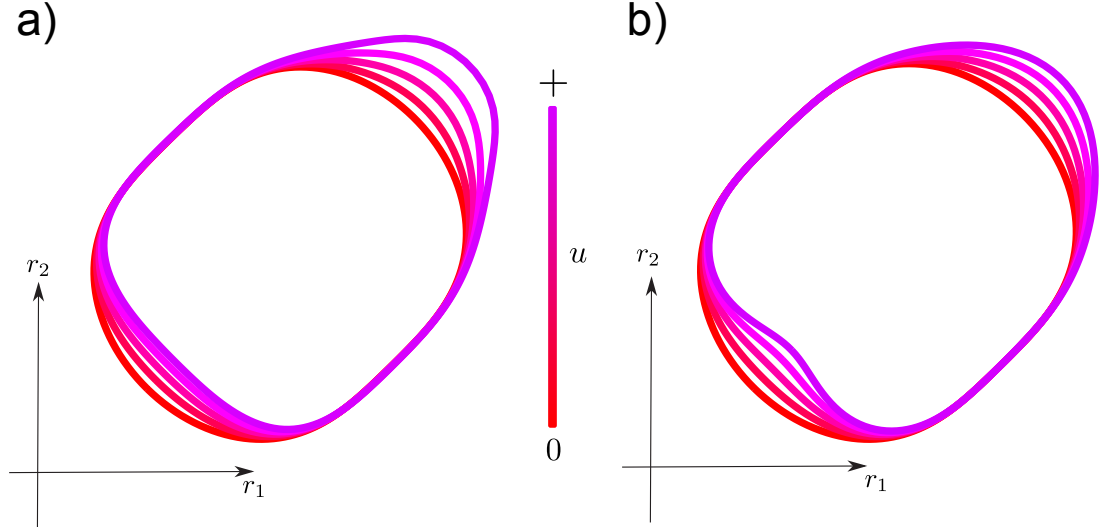
In general, it takes a handful of parameters to describe gaits, and the outcome of gaits is of relatively small size (i.e., SE(2) has only 3 degrees of freedom). A higher level control is useful in simplifying the control complexity. Previously discussed gait modulation naturally enables a high level control over the robot movement despite its complex locomotor behavior.

In an online, within-cycle type of control scenario, gait switching/modulation can happen at phases where gaits are not collocated. Large-amplitude “jumps” can cause the robot to stop the current motion, reposition its joints, and then finally follow the new gait, which in general can badly affect the overall smoothness and efficiency of the system motion. Limiting the bandwidth of the modulation rate can prevent jumps from one gait to another. This limitation acts intuitively like a dynamically adjusted interpolation between current and target gaits.

In an online control situation, the commanded shape velocity  $\dot{r}$  is constructed as

$$\dot{r}(t) = \dot{r}^* + \dot{u}d(t) + u\dot{d}(t), \tag{3.14}$$

where  $\dot{r}^*$  denotes the shape velocity along the current gait,  $u$  denotes the control variable providing steer-ability, and  $d(t)$  denotes the difference between the current gait and target gait at time  $t$ . Because drag-dominated systems are most efficient



**Figure 3-2.** An illustration of gait modulations generated with (a) gradient iteration and (b) first-order approximation at the nominal gait. When the modulation variable  $u = 0$ , the gait (red) is the x-direction optimal gait (i.e., the same as the one shown in Fig. 3-1(b)). As  $u$  increase in value, the nominal gait is modulated toward the top-right corner for additional  $\theta$  movement.

when their power dissipation is held constant, we want to limit  $\dot{u}$  such that it does not significantly impact the metric-weighted speed of the system,

$$\|\dot{r}\|_{\mathcal{M}} \approx \|\dot{r}^* + u\dot{d}(t)\|_{\mathcal{M}}. \quad (3.15)$$


In the demonstration below, we implement this restriction enforcing

$$\|u\dot{d}(t)\|_{\mathcal{M}} \leq k\|\dot{r}\|_{\mathcal{M}}, \quad 0 < k \ll 1. \quad (3.16)$$

### 3.4 Results

In this section, we provide a parallel comparison among different approaches in generating gait modulations using the example simulated system three-link swimmer mentioned in 3.2. Also, the following results are generated under the circumstance where an optimal forward gait is modulated for turning.

Following the methods mentioned in 3.3.1, a number of snapshots from the solution of (3.11),  $p(u)$ , are shown in Fig. 3-2(a). As expected, as the modulation variable  $u$




**Figure 3-3.** An overview of the performance of the gait libraries from the global approach. (a) Gaits generated with the global approach, where  $\alpha = 1, 0$  corresponds to those optimal gaits shown in Fig. 3-1(b-c). (b) Displacement in  $(x, y, \theta)$  per cycle of the gait library. Note that the discontinuity appears around the region  $\alpha = 0.4$  where the optimizer decides to qualitatively change the gait. (c) Pareto front, generated from (3.13), on the efficiency trade-off map. The discontinuity appears here as a large gap where the optimizer was not able to reliably converge to data points in between, likely due to a bifurcation in the objective function as a function of  $\alpha$ .

increases, the gradient pulls the gaits toward the top-right corner, where the optimal rotation gait is located (see Fig. 3-1(b)).

As discussed in 3.3.1, under waypoint parametrization, the gradient is given by local definitions along the gait. Consequently, part of the gait can be trapped in some local critical points, where the gradient diminishes for some waypoints, even though other parts of the gait are advancing with the gradient. In Fig. 3-2(a), the gradient along the two sides (top-left and bottom-right) started as zeros, leaving no changes across modulations. Similarly, as the bottom-left side of the gait gets close to a local critical point, this side converges and is not able to proceed onward.

The first-order approximation renders the gait modulation linearly as expected. The effect on the bottom-left side differs from the one mentioned above; because the gradient is non-zero at the start, that side will keep modulating even though it causes additional unnecessary path-length costs by dimpling inwards.

There is a significant gap in the middle of the Pareto Front (see Fig. 3-3) around  $\alpha = 0.4$ , where the optimizer is giving up a large amount of x-direction efficiency in



**Figure 3-4.** An efficiency comparison between the gait modulations generated by local and global approaches. The performance of the gradient-based methods is almost as good as the global optimizer locally at the x-optimum, and slowly drops off.

return of  $\theta$ -direction efficiency. We conjecture that the formation of such a bifurcation region is due to the balancing of two local optima in  $x$  and  $\theta$  efficiencies, similar to those observed in the gradient iteration method. Around the region near  $\alpha = 0.4$ , the two local optima yield close values. As  $J_{total}$  is altered by  $\alpha$ , the global optimum has a sudden jump, both in the appearance of the gait and the motion output. Further investigation is needed to make a conclusion here.

Intuitively such bifurcation is reasonable, where there are certain limits to the extent that modulating a gait is meaningful, once beyond that limit, then it is time to switch to some other qualitatively different gait.

Fig. 3-4(a) shows a zoomed-in view of the three methods on the trade-off map between  $x$  and  $\theta$  efficiencies, focusing at the region close the x-direction optimum (bottom right corner of Fig. 3-3). The efficiency of gradient-based methods slowly drop off as they move towards turning, part of the reason being segments of the gait are trapped at local critical points.

## 3.5 Demonstration

We implemented our gait modulation generation along with an additional command filter control described in 3.3.3. The following demonstration is made on the same example system as in previous sections, the gait libraries are generated by the gradient iteration methods.

We chose turning curvature as an intuitive control input. For the gait modulations, average turning curvature, which is a function of the  $x$  and  $\theta$  displacements induced by executing one full gait cycle, was used to map desired curvature to the gait modulation variable  $u$ . Under an arbitrary command (i.e., to perform a slight turn during forward motion), the controller limits the rate of modulation variable,  $\dot{u}$ , based on the differences between current and target shape as discussed in (3.16), where  $k$  is set to 0.1. The gait trajectory in Fig. 3-5 is as expected, when at a phase where current gait and target gait are not collocated, the command filter will dynamically interpolate between the two gaits, preventing sudden jumps in the shape space.

## 3.6 Discussion and Conclusion

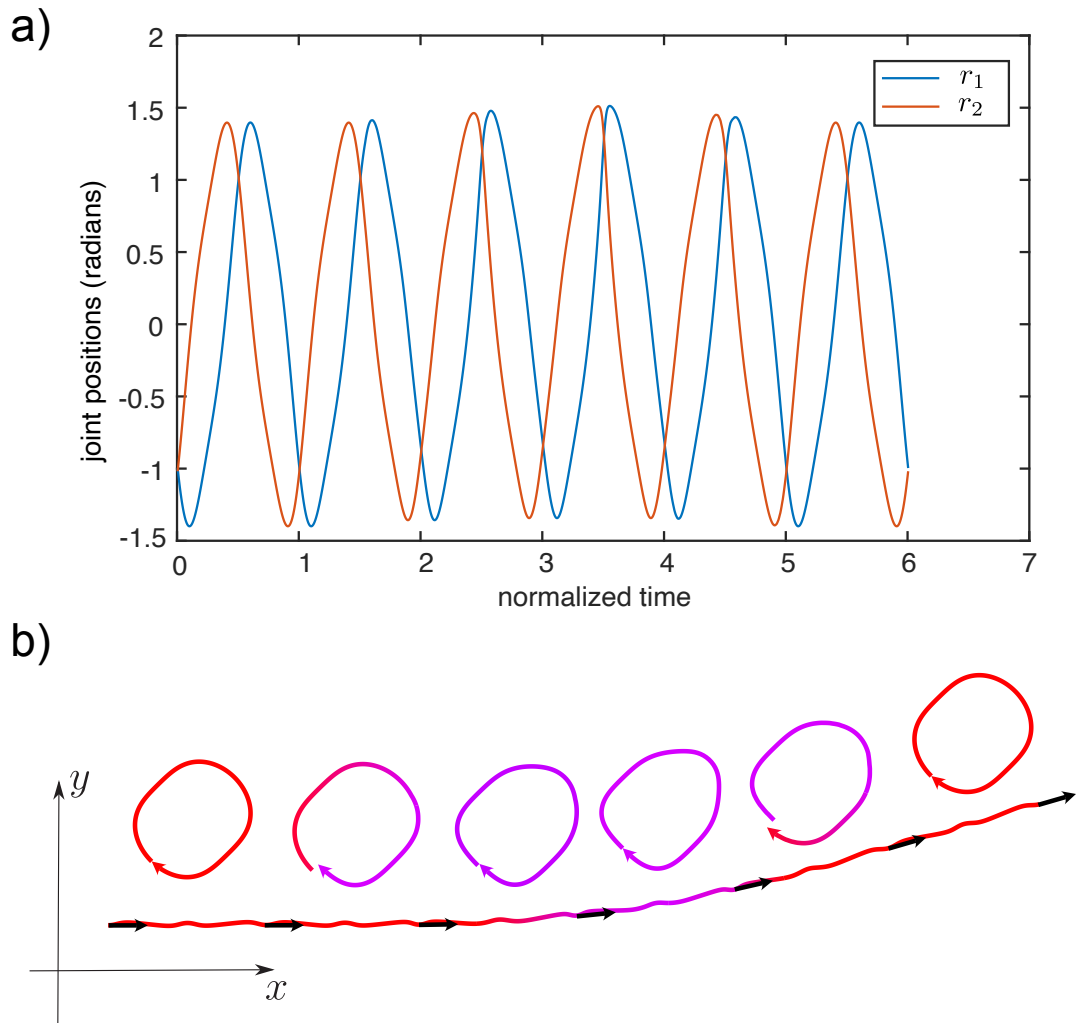
In this chapter, we incorporate insights from geometric gait optimization to construct gait modulations that enable “controllable maneuverability”. The local approach we present involves starting with a forward-optimized gait, and then inducing gait parameter modulations that push the gait in the gradient direction of a turning

objective function; this local strategy provides an effective turning control system up to moderate turns, according to our numerical investigation of a simplified swimmer. The global approach aims to generate a continuous family of optimal gaits that satisfy Pareto optimality—trading off forward movement and turning. We also discussed the advantages and limitations between local and global approaches in generating gait modulations.

Finally, we demonstrated a common use case of gait modulations, where a high-level continuous steering control is possible while largely maintaining a nominal forward gait.

Some early works [122] are successful in building hydrogel-based origami structures in which they are able to achieve a variety of shape changes including bending, elongating, twisting, and buckling. From a roboticist point of view, these shape changes can be utilized as actuation for soft robots. On the other hand, there is no obvious way to build low-dimensional, first-principles models for the locomotor dynamics—making our interest in using local maneuvering based on data-driven models [68] particularly salient. Ultimately, we are interested in putting together micro-scaled, hydrogel-based soft robots that locomote and perform a wide range of dynamic, maneuverable locomotor tasks based on environmental information.

For the high-deformation structures, a key focus of future work will be understanding how to access the steering modulations when only some of the shape variables are directly controlled. This work will combine the ideas outlined in this chapter with the geometric mechanics of semi-passive locomotion we explored in [123].



**Figure 3-5.** A demonstration of gait modulation in motion. (a) Shape trajectory, as quasi-periodic functions, of the robot over a normalized time (i.e. one full lap in the shape space per normalized time). (b) The system's spatial trajectory by executing the above shapes. Black arrows denotes the system's  $x$ -axis position at the beginning of each lap, and the color coded trajectory shows approximately (not to scale) the average turning curvature is currently being performed. The shape space trajectory of each lap is also shown above the trajectory as a simple visualization.



## Chapter 4

# Data-Driven Geometric Modeling of Systems with Low-Bandwidth Actuators

This chapter was submitted for publication and was made available on arXiv entitled *A Data-Driven Approach to Geometric Modeling of Systems with Low-Bandwidth Actuator Dynamics* authored by Siming Deng, Junning Liu, Bibekananda Datta, Aishwarya Pantula, David H. Gracias, Thao D. Nguyen, Brian A. Bittner, Noah J. Cowan. We re-used the article in this chapter with permission from all the authors.


### 4.1 Motivation

Many conventional robots rely predominantly on rigid, fully actuated mechanisms. While these robots maintain superior force and precision compared to natural organisms, these rigid machines usually struggle in tasks that involve interacting safely with humans, handling deformable objects, and operating in unstructured environments [124]. Designs from nature have inspired the development of compliant mechanisms in robotics, enabling new capabilities [14, 125–131]. The emergence of such soft components in robotic platforms has provided new avenues for improved adaptability, safety, cost, and energy efficiency. On the other hand, these new mechanisms introduce new challenges in modeling and control. The compliant nature of these soft

components greatly increases the internal degrees of freedom as well as the degrees of underactuation, which makes it hard to obtain precise control of each shape element simultaneously. One potential approach is to utilize variations in passive actuation responses to stimuli across different parts of the body as a means of locomotion. This strategy capitalizes on the distinct temporal dynamics among subsystems or constituents.

In addition to compliance, underactuation, and low-bandwidth dynamics, there are other challenges introduced by soft robotic systems that complicate modeling efforts. For example, boundary conditions, surface interactions, and nonlinear material properties make it difficult to derive parsimonious models from physical principles. Using a top-down, data-driven approach, this work investigates the aforementioned passive responses within a systematic framework for locomotion control. We focus on systems with low-bandwidth shape changes in response to a single actuator input. Modeling actuator dynamics and its effects on the system can streamline engineering efforts to design and control soft robots, maximizing their capabilities with less exploratory or exhaustive experimentation. In many cases, a global model of system dynamics may be impractical to obtain, especially for custom machines, machines made through imprecise fabrication techniques, or systems deployed in poorly characterized environments (e.g. a patient’s body). A key insight of this work is that in these scenarios, we can avoid reliance on a global model or the sample inefficiency of reinforcement learning schemes by systematically achieving function through iterative modeling and refinement of local behaviors.

Geometric mechanics provides a framework through which top-down insights permit novel motion planning approaches to dissipative systems with symmetry, including analytical computation of optimal gaits [58, 120]. A core premise of this work is that complex locomotor mechanics can be rewritten in a kinematic form, owing to the assumption that dissipative forces dominate the physical interaction between the body



**Figure 4-1.** Experimental screenshots of thermal cycling of two-segment robots with a flexible linker at the end of a cooling half-cycle (top), and heating half-cycle (bottom). One thermal cycle comprises a heating half-cycle and a cooling half-cycle. The robot displaces 4.4% body length at the end of the cycle.

and environment [42]. The same framework has been instrumental in understanding cyclic locomotion in nature [73, 132]. Bittner et al. [68] presented a data-driven approach to construct a local model in the neighborhood of the observed limit cycle, using data points from stochastically perturbed, repeated behaviors. More recent work [71] extended this data-driven approach to shape-underactuated systems, which have high bandwidth control available only to a subset of the shape space. The ability to build local models provides the opportunity to sample candidate gaits offline for sample efficient hardware-in-the-loop optimization.

In contrast to hard robots that can readily be powered by multiple, independent electromagnetic modules with rapid (i.e., high-bandwidth) responses, soft robots often must be stimulated by an external signal, typically a single stimulus such as

pH, a specific biomolecule, light, or temperature. The impact of this signal must play out through complex kinetics (i.e., *low-bandwidth control*). For this class of systems, shape deformation is not instantaneously coupled to the control signal; rather, there is a temporal lag in the excitation of actuator dynamics following the control input. In our prior work [75] (see Fig. 4-1). In this prior work, we conceptualized and built a thermo-responsive hydrogel crawler. Although stimulus-responsive shape changes for hydrogels are ubiquitous in literature [133–135], the design of our robot exploited the swelling and shrinking induced bending mechanism, morphological asymmetry, and asymmetry in friction force in response to the change in surrounding temperature to achieve the locomotion. In this crawler, there are three distinct segments: a suspended linker segment connects two end bilayer segments comprising active poly(*N*-isopropylacrylamide) (pNIPAM) and passive polyacrylamide (pAAM) layers with different morphologies. Asymmetry in friction forces, caused by morphological asymmetry, between the two bilayer segments at low and high temperatures allowed the robot to change its anchor during a temperature cycle to move unidirectionally. Additionally, we hypothesize that the distinct swelling rates of these bilayers create such asymmetric ground reaction forces that can be exploited for locomotion. Utilizing the asymmetric response time among segments, this robot can locomote with a single cyclic input—temperature cycle. Alongside the fabrication of this physical crawler, we also developed a Finite Element Analysis (FEA) model [75] to simulate the response and investigate the deformation mechanisms, which we use to generate the results in this chapter.

Our core contribution, presented in Section 4.3, is the data-driven modeling of a ubiquitous class of underactuated systems, where the shape dynamics are subject to a band-limited control. This contrasts with prior work [71], which required the application-limiting assumption that at least one element can be accessed by high bandwidth control. Our work enables the application of our framework for efficient

behavior optimization and enhancement across a wide spectrum of previously unexplored soft robotic systems, such as hydrogel crawlers [71]. In 4.4, we demonstrate our methods on a well-known, analytically tractable system, modified to include low-bandwidth actuation of its shape parameters. Finally, in Section 4.5 we test our methods on a high-dimensional, finite element model of our previously published hydrogel robot [75]. In both examples, we show how the actuator dynamics can be simultaneously modeled with the body movements, enabling a data-driven modeling architecture for a broader class of soft or underactuated systems. Further, we use these examples to numerically optimize a parameterized input signal for certain objectives using an iterative parameter optimization and model refinement approach.

## 4.2 Background

### 4.2.1 Geometric Locomotion Model

Geometric mechanics [42, 48, 101] provides a framework for locomotion based on exploiting symmetry. In this framework, a distinction is made between the internal configuration (shape) of a locomotion system and its position and orientation (group) in a spatially fixed reference frame. Central to this framework is the idea of group invariance of the dynamics [47]: a shape change that moves the system in a certain way—in the system body frame—will do the same at any position and orientation in the environment, invariant to the absolute position and orientation.

Here, we consider a subclass of such group-invariant dynamical systems that are dominated by dissipative forces as caused by many types of isotropic friction [102]; in such dissipation-dominated systems, the equations of motion can be kinematically reduced such that the body velocity is expressed as a shape-dependent linear mapping of shape velocity. In this case, the kinematic equation is

$$\left(g^{-1}\dot{g}\right)^\vee = \xi = -\mathbf{A}(r)\dot{r}, \quad (4.1)$$

where  $\xi$  is the group velocity in its body frame<sup>1</sup>,  $r$  denotes the system shape, and  $\mathbf{A}(r)$  is called the *local connection*. Here the matrix  $\mathbf{A}(r)$  is a function of shape  $r$  and acts analogously to a Jacobian in which it relates the system’s shape velocities to body velocities. A spatial trajectory of the system body frame can be calculated by integrating (4.1) with respect to a fixed reference frame.

For systems within the scope of (4.1), the local connection can be analytically derived from a set of Pfaffian constraints on the system’s shape and body velocities. A global model can be empirically estimated by exhaustively sampling the system shape space and its tangent bundle (the collection of shape velocities available at each point in the shape space) [60]. However, such global models are often difficult to obtain for animals or underactuated systems because of the challenges in sampling this space with sufficient density.

## 4.2.2 Data-driven Modeling

Bittner et al. [68] developed a data-driven approach to geometric modeling and optimization, which was later extended [71] to be applied to systems with high bandwidth control in only a subset of the shape variables. This approach allows local estimation of a connection in the neighborhood of a limit cycle with far fewer samples than required to train a global model, making it practical for *in-situ* system identification, especially for systems with high dimensional shape spaces.

In this approach, system shape data (in the form of a regularly sampled time series) are fit to an oscillator such that each data point is assigned a phase value [136]. A zero-phase-lag Butterworth smoothing filter is applied before finite differencing to obtain time derivatives of both shape and position. Then, a local Taylor approximation of the connection can be computed via linear regression across data points within

---

<sup>1</sup>Here  $(\cdot)^\vee$  is an isomorphism that maps velocities from Lie algebra form to vector form, and its inverse is denoted as  $\widehat{(\cdot)}$ . In a  $SE(2)$ ,  $(\cdot)^\vee : se(2) \rightarrow \mathbb{R}^3$ , and  $\widehat{(\cdot)} : \mathbb{R}^3 \rightarrow se(2)$

phase windows. A Fourier series is then fit to these local regression coefficients to build a model that is supported at any queried phase.

We detail the process by which we estimate a linearized model within each phase window. Data-driven Floquet analysis techniques extract information from the observed oscillator data and assign each sample point an estimated phase [136, 137]. The observed shape samples are then phase-averaged and fitted to a Fourier series to obtain a limit cycle, denoted as  $\theta_r(\cdot)$ . The perturbed trajectory, relative to the limit cycle, is denoted as  $\delta_r := r - \theta_r$ . The first-order Taylor approximation of the local connection in each local phase window can be constructed as

$$\mathbf{A}_k(r) \approx \mathbf{A}_k(\theta_r) + \delta_r^T \frac{\partial \mathbf{A}_k}{\partial r}, \quad (4.2)$$

where  $\mathbf{A}_k(r)$  is the  $k^{\text{th}}$  row of the local connection, which is a vector of the same dimension as shape perturbation  $\delta_r$ .

All samples are grouped into neighborhoods by their estimated phase values, and a local model is fitted in each phase window. In the  $m^{\text{th}}$  phase window, the averaged shape is assumed to be a constant  $\theta_r^m$ . The first-order Taylor approximation of the local connection matrix  $\mathbf{A}(r)$  in this phase window can be fitted by solving the following Generalized Linear Model (GLM):

$$\xi_k^{(n)} \sim \mathbf{C}_k + \mathbf{B}_k \delta_r^{(n)} + \mathbf{A}_k(\theta_r) \dot{\delta}_r^{(n)} + \frac{\partial \mathbf{A}_k}{\partial r} \delta_r^{(n)} \dot{\delta}_r^{(n)}. \quad (4.3)$$

Here,  $\xi_k^{(n)}$  corresponds to the  $k^{\text{th}}$  coordinate of the  $n^{\text{th}}$  sampled body velocity  $\xi^{(n)}$ , and  $\delta_r^{(n)} := r^{(n)} - \theta_r^m$ ,  $\dot{\delta}_r^{(n)} := \dot{r}^{(n)} - \dot{\theta}_r^m$  are the shape and shape velocity perturbation samples defined in the local region indexed by  $m$ . Regressor  $\mathbf{C}_k := \mathbf{A}_k(\theta_r) \dot{\theta}_r$  describes the average behavior in the neighborhood of  $\theta_r^m$ .  $\mathbf{B}_k := \dot{\theta}_r^T \frac{\partial \mathbf{A}_k}{\partial r}$  and  $\mathbf{A}_k$  are the terms that respectively relate the effects of shape and shape velocity offsets from the limit cycle.  $\frac{\partial \mathbf{A}_k}{\partial r}$  is the cross term that incorporates the interaction between  $\delta_r$  and  $\dot{\delta}_r$ . Note that this is a local estimate in the  $m^{\text{th}}$  phase window. This local approximation is

repeated for all separate groups of data points, after which a Fourier series model is used to guarantee a smooth transition among the fitted matrices.

## 4.3 Methods

### 4.3.1 Low-bandwidth Shape Control

In this chapter, we consider systems whose locomotion can be characterized by (4.1) while only having access to low-bandwidth control over  $r$ . In particular, we assumed the dynamics on  $r$  to take the general form of

$$\dot{r} = f(r, u), \quad (4.4)$$

where the system shape velocity  $\dot{r}$  is a function of its shape  $r$  and an input  $u$ .

First, we extracted a phase-averaged gait cycle  $(\theta_r, \theta_u)$  from the general input [136]. Denoting the perturbation from phase-averaged shape and control as  $\delta_r := r - \theta_r$ ,  $\delta_u := u - \theta_u$ , the local first-order Taylor approximation of the actuation dynamics can be written in the following form:

$$f(r, u) \approx f(\theta_r, \theta_u) + \frac{\partial f}{\partial r}(\theta_r, \theta_u)\delta_r + \frac{\partial f}{\partial u}(\theta_r, \theta_u)\delta_u \quad (4.5)$$

We then fit the data to the above first-order approximation by solving the following Generalized Linear Model,

$$\dot{\delta}_r^{(n)} \sim \mathbf{D} + \mathbf{E}_r \delta_r^{(n)} + \mathbf{E}_u \delta_u^{(n)}, \quad (4.6)$$

where  $D$  is the average shape velocities of the observed data in the local phase window, and  $(\mathbf{E}_r, \mathbf{E}_u)$  are the terms that describe how shape and input offsets respectively modify the average behavior.  $\delta_r^{(n)} := r^{(n)} - \theta_r^m$ ,  $\delta_u^{(n)} := u^{(n)} - \theta_u^m$  are the shape and input perturbations defined in the  $m^{\text{th}}$  local phase window, where  $(\theta_r^m, \theta_u^m)$  are the mean values of shape and input.

The estimation of the local connection can be done separately from the actuator dynamics, hence this part remains identical as in 4.2.2. We repeat the same procedure



for all discrete phase windows and use a Fourier series to smoothly connect all local models.

The fitted models from (4.6) and (4.3) can be used in series to make predictions of the system shape and position trajectories given the input signal. First, the input signal  $u(t)$  is transformed into the phase coordinate using the fitted phase map. The initial shape is assumed to be on the limit cycle ( $\delta_r = 0$ ) as the same phase value of the initial input  $u(t_0)$ . At each discrete time  $t_i$ , the shape velocity perturbation  $\dot{\delta}_r(t_i)$  is predicted using the actuator model (4.6) given the current shape perturbation  $\delta_r(t_i)$  and the input perturbation  $\delta_u(t_i)$ .  $\delta_r(t_i)$  is then integrated by Euler's method to obtain the predicted shape at the next time step  $\delta_r(t_{i+1})$ . The predicted shape  $\delta_r(t_i)$  is then used to predict the body velocity  $\xi(t_i)$  using the body velocity model (4.3). The predicted body frame position at the next time step  $g(t_{i+1})$  is then integrated using  $\xi(t_i)$ .

When building an actuator model, the system shape trajectory is the integrated estimation of the shape velocity predictions. Simultaneously, it also appears as the input to the local connection in the locomotion model. As in prior work, the shape and body motion models are predicted in separate stages. Note that in the process of simulating a system spatial trajectory from a general input signal, the two integration steps of each model evolve in series. We start with knowledge of the initial system shape  $r(t = 0)$  and the control input  $u(t)$ . Then we can numerically solve (4.4) and (4.1) together using the fitted regression models, (4.6) and (4.3).

We apply the model improvement metric described in [71], comparing our first-order regression model predictions to the phase-averaged baseline model predictions,

$$\Gamma_\chi = 1 - \frac{\sum_{n=1}^{\mathcal{N}} \|\chi_D^{(n)} - \chi^{(n)}\|}{\sum_{n=1}^{\mathcal{N}} \|\chi_T^{(n)} - \chi^{(n)}\|}. \quad (4.7)$$

This improvement metric is defined as one minus the relative error of the data-driven prediction  $\chi_D$  with respect to the baseline prediction  $\chi_T$  over  $\mathcal{N}$  samples of body

velocity and shape velocity  $\chi = \{\xi, \dot{r}\}$ .  $\Gamma_\chi \leq 0$  means the data-driven prediction is no better than the phase-averaged prediction, and  $0 < \Gamma_\chi \leq 1$  means that our model can make better predictions than the baseline model, up to perfect reconstruction of the ground truth at  $\Gamma_\chi = 1$ .

The baseline phase model corresponds to the zeroth-order model, generating predictions solely reliant on current phase information. This metric holds significance in assessing the extent to which our data-driven first-order model outperforms the baseline model within the local region. A substantial improvement metric implies that the data-driven model more accurately approximates the ground truth compared to the baseline model within the local region, making it more suitable for local optimization.

### 4.3.2 Optimizing Behaviors and Iterative Model Refinement

Once an initial model is obtained, we can make predictions on the system position trajectories  $g(t)$  by a general control input  $u(t)$ . Using finite differencing, we can estimate the gradient of displacement per cycle with respect to the control parameters around the observed data. We can then utilize the estimated gradient and Hessian to numerically optimize the control parameters for certain behaviors of the system (e.g. maximizing the displacement per cycle).

The expense of data collection<sup>1</sup> incentivizes our focus on sample-efficient optimization schemes. Given an input parameterization, we sparsely sample data in the full range of the input space and build a rough model. According to this rough model, we numerically optimized the input parameters for certain behavioral objectives. Then we zoom into the region around the optimized parameter and re-sample points in this local area. The model built with sample points in a smaller region will be more localized and accurate over that confined domain. We iterate between these two processes—optimization and model refinement—so that in the end it converges to a local optimum in the control space. A global optimum is not guaranteed. In Sections

4.4 and 4.5, we demonstrate the methods using a sample objective where we maximize the displacement per cycle and penalize the cycle time.

## 4.4 Illustrative Example: Purcell Swimmer with Low-bandwidth Actuation

Before demonstrating the method on data, it is helpful to interrogate a simple analytical model, such as a three-link Purcell swimmer. We modify the model to include low-bandwidth actuation, inspired by the low-bandwidth actuation of our hydrogel robot:


$$\dot{r}_i = c_i(r_i^s(T) - r_i), \quad c_i > 0, i = 1, 2, \quad (4.8)$$

where  $r_i$  is the  $i^{\text{th}}$  shape variable (joint positions),  $r_i^s(T)$  denotes each joint’s steady state equilibrium given temperature, and  $c_i$  is the converging rate of each joint towards its steady state equilibrium. Specifically, the steady-state equilibrium  $r_i^s(T)$  is assigned to be a linear function of a one-dimensional input signal, temperature  $T$ . We have a bound on temperature that puts limits on the swimmer’s joint angles. The resulting shape trajectory exhibits hysteresis which is often observed in low-bandwidth control systems (see Fig. 4-2).

Assuming different constants  $c_i$  on the two joints, the shape variables will exhibit a gait where both joints are not synchronized under a repeating temperature cycle, see Fig. 4-3. Although both joints are controlled by the same temperature input, the phase lag between the two joints breaks the symmetry of joint synchronization, making the gait enclose a nonzero area in the shape space, which is critical for locomotion in viscous swimming domains from the scallop theorem [104].

---

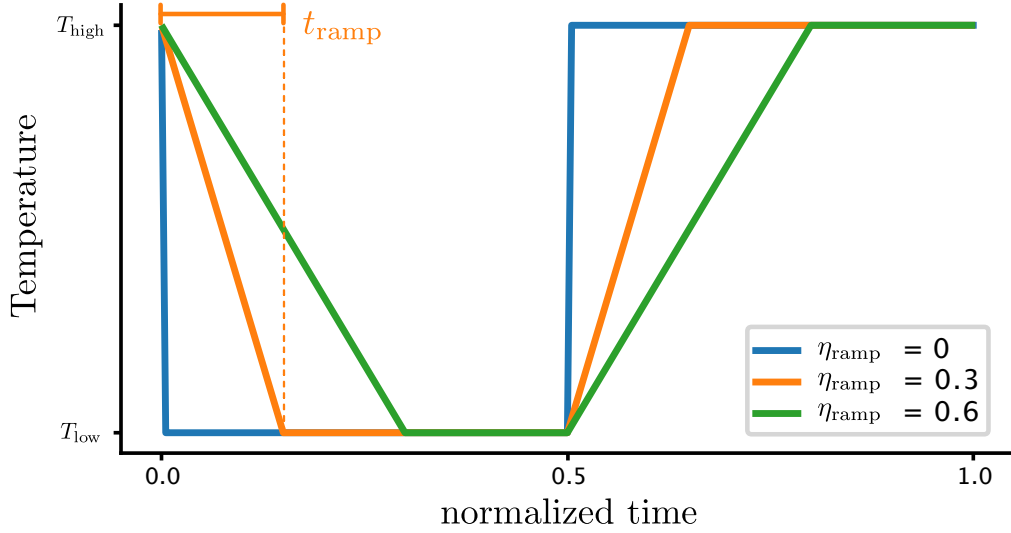
<sup>1</sup>In practice, a typical temperature cycle for our hydrogel crawler takes approximately 6 hours because of the slow actuation kinetics of the material. The FEA simulation is computationally expensive as well; running a 10-cycle simulation on a well-equipped desktop computer takes about 2 days. Both facts make data collection for such systems expensive, thus data efficiency is crucial.



**Figure 4-2.** Plotted on the shape space, phase-lag-induced shape trajectories exhibit the hysteresis often observed in low-bandwidth control systems. The greater the difference in the reaching rate  $c_i$ , the larger the phase lag is. This behavior emerges by waiting for a sufficient amount of time for the two joints to both reach steady-state equilibrium and then exciting both joints by a step change in input.

#### 4.4.1 Input Generation

Our parameterization on the control signal is concise while maintaining the ability to alter important features of the temperature profile. Here, we used 4 parameters to describe the temperature cycle: a low-point temperature  $T_{\text{low}}$ , a high-point temperature  $T_{\text{high}}$ , time-span per cycle  $t_{\text{cycle}}$ , and the portion of the half period to ramp between high and low temperatures  $\eta_{\text{ramp}} = 2t_{\text{ramp}}/t_{\text{cycle}}$ , where  $t_{\text{ramp}}$  is the time to ramp between high and low temperatures, see Fig. 4-3. Performing multiple cycles of these parameterized temperature cycles, the shape trajectory forms a stable orbit under



**Figure 4-3.** Example temperature cycles (our low-bandwidth control input), normalized for period and temperature, where  $t_{\text{ramp}} = \eta_{\text{ramp}} \cdot t_{\text{cycle}}/2$  is the time used to ramp input to the goal temperature. Since the ramp time should always be smaller than half of the period, the use of the ramp time ratio avoids additional constraints other than a bounding box. When  $\eta_{\text{ramp}}$  is small, the temperature profile is more square-wave-like, and the resulting shape trajectory is more likely to have a larger phase lag. Consequently, for longer ramp times, the shape trajectory is more likely to have minimal phase lag because both joints are in equilibrium for the duration of temperature change.

periodic forcing. We then perturbed the forcing parameters across cycles, which resulted in what can be seen as a "tube" around the orbit as shown in Fig. 4-5.

## 4.5 Main Application: Hydrogel Crawler

### 4.5.1 Actuator Dynamics

Bilayers and other multi-material structures are useful in creating interesting modes of shape changes like bending. Typical swelling-driven bilayer bending dynamics are similar to the form of an exponential low-pass filter as shown in Fig. 4-2. Specifically, the geometry of the bilayer (e.g., layer thickness ratio and material properties) can affect the steady state equilibrium of the shape variables as well as the rate of reaching equilibrium.

### 4.5.2 Hydrogel Crawler

Thermo-responsive hydrogel crawlers in [75], capable of swelling and shrinking, utilize geometric asymmetry, leading to asymmetry in friction forces, to generate net motion under temperature cycles. We utilized the same 2D finite element model in Abaqus Unified FEA [138] to produce a time-dependent  $x - y$  coordinate along the contour of the robot to calculate the area (2-D volume) of the actuated segments. The data is then parameterized into the shape variable  $r$ .

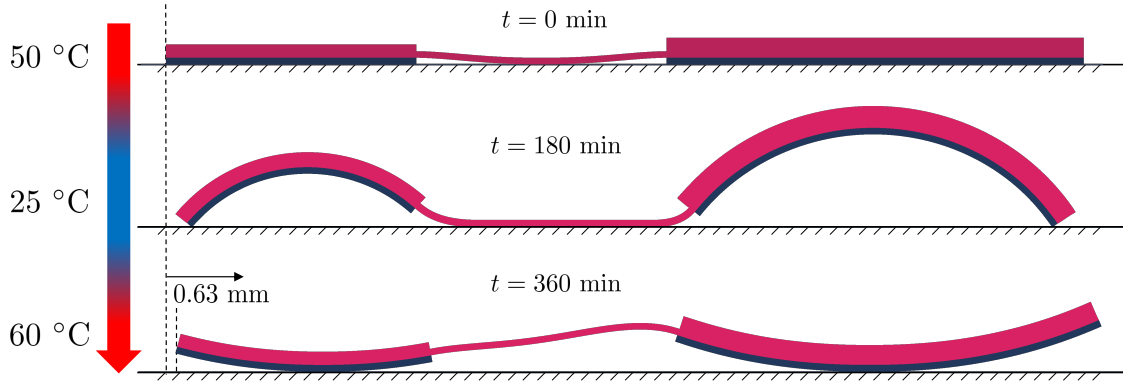
Here we assume the actuation dynamics in a general (nonlinear) form (4.4), without any specific structure on it, namely

$$\dot{r} = f(r, T), \tag{4.9}$$

where the input is assumed to be the temperature  $T$ .

### 4.5.3 Finite Element Model

Briefly, our finite element model, based on chemo-mechanics described in [139], solves coupled diffusion-deformation equations for hydrogel undergoing temperature-driven swelling and shrinking. We used Neo-Hookean and Flory-Huggins potentials to describe the entropic elastic behavior of the polymer network and the mixing of polymer-solvent, respectively. The swelling of pNIPAM caused by the lower critical solution temperature transition (LCST) was modeled by assuming a sigmoidal function for the temperature dependence of the Flory-Huggins interaction parameter. We also assumed that the diffusivity of the water through pNIPAM also increased sigmoidally with temperature across the LCST, which caused the characteristic time of deswelling to be significantly faster than the characteristic time of swelling. Given the relatively long actuation time, the environment temperature can be simply assumed to be evenly distributed. We also considered a combined effect of gravity and buoyancy by prescribing a net body force on the hydrogel structure. Our material model included a total of 10 parameters which



**Figure 4-4.** A representative finite element analysis showing unidirectional motion towards the larger bilayer of the thermo-responsive hydrogel crawler subjected to 3 hours of ramped cooling and 3 hours of ramped heating cycle.

Parameter	Value
Differential density between polymer and water, $\Delta\rho$	100 kg/m <sup>3</sup>
Shear modulus of pNIAPM, $G_{\text{pNIAPM}}$	2.32 kPa
Shear moduli ratio of the gel, $\frac{G_{\text{pAAAM}}}{G_{\text{pNIAPM}}}$	15.36
Curing temperature, $T_{\text{cure}}$	323 K
Flory-Huggins parameter at high temperature, $\chi_H$	0.8031
Flory-Huggins parameter at low temperature, $\chi_L$	1.0483
Transition temperature, $T_{\text{trans}}$	318 K
Range of transition temperature, $\Delta$	8.0
Diffusivity at low temperature, $D_L$	10 <sup>-11</sup> m <sup>2</sup> /s
Diffusivity at high temperature, $D_H$	10 <sup>-10</sup> m <sup>2</sup> /s

**Table 4-I.** List of materials properties used in the simulation

were either directly determined from experiments or calibrated against experiments using finite element analysis. A list of the parameters and their values are listed in Table 4-I. In addition, we assumed a rigid frictional surface with a friction coefficient,  $\mu_k = 0.1$ , underneath the robot to facilitate friction-driven locomotion induced by geometric asymmetry. Further details of the finite element simulation can be found in [75].

#### 4.5.4 Input Generation

Here we used a similar parameterization for the input temperature signal as that in 4.4.1. However, for thermo-responsive hydrogels, the kinetics of the swelling and shrinking vary distinctively. We, therefore, separated the input cycle into two independent parts, cooling and heating. Thus, the dimension of the input parameter space increases to six, low temperature  $T_{\text{low}}$ , ramped cooling time span  $t_{\text{cool}}$ , cooling ramp time ratio  $\eta_{\text{cool}}$ , high temperature  $T_{\text{high}}$ , ramped heating time span  $t_{\text{heat}}$ , and heating ramp time ratio  $\eta_{\text{heat}}$ . The ramp time ratios are defined as the ratio of the corresponding ramp time to the cooling or heating time span, i.e.,  $\eta_{\text{cool}} = t_{\{\text{ramp,cool}\}}/t_{\text{cool}}$ . The allowable range of each parameter is determined by material properties and characteristic diffusion time and was validated by a parametric study using FEA. Specifically, the ramped cooling and heating timespans are determined by scaling swelling and shrinking characteristic time of the hydrogel, both temperature ranges are specified by 4% equilibrium strain span of the material. The ramp ratios are ideally in the range of  $[0, 1]$ , but small ramp ratios mean very large rates of temperature change, which is impractical and often causes numerical stability and convergence issues in FEA simulation because of the excessive deformation of the finite element mesh in a short period of time. Thus in the implementations, we raised the lower bounds to  $\frac{1}{32}$ . The calculated full input parameter ranges are shown in Table 4-II. The noisy input signal is generated by sampling from a uniform distribution in the parameter space. The input parameters are then used to generate noisy temperature cycles for FEA. To avoid numerical issues, instead of running hundreds of thermal cycles at a time, each FEA simulation comprised 10 thermal cycles. At each iteration, we ran 10 simulations resulting in 100 cycles of input data for our data-driven model.

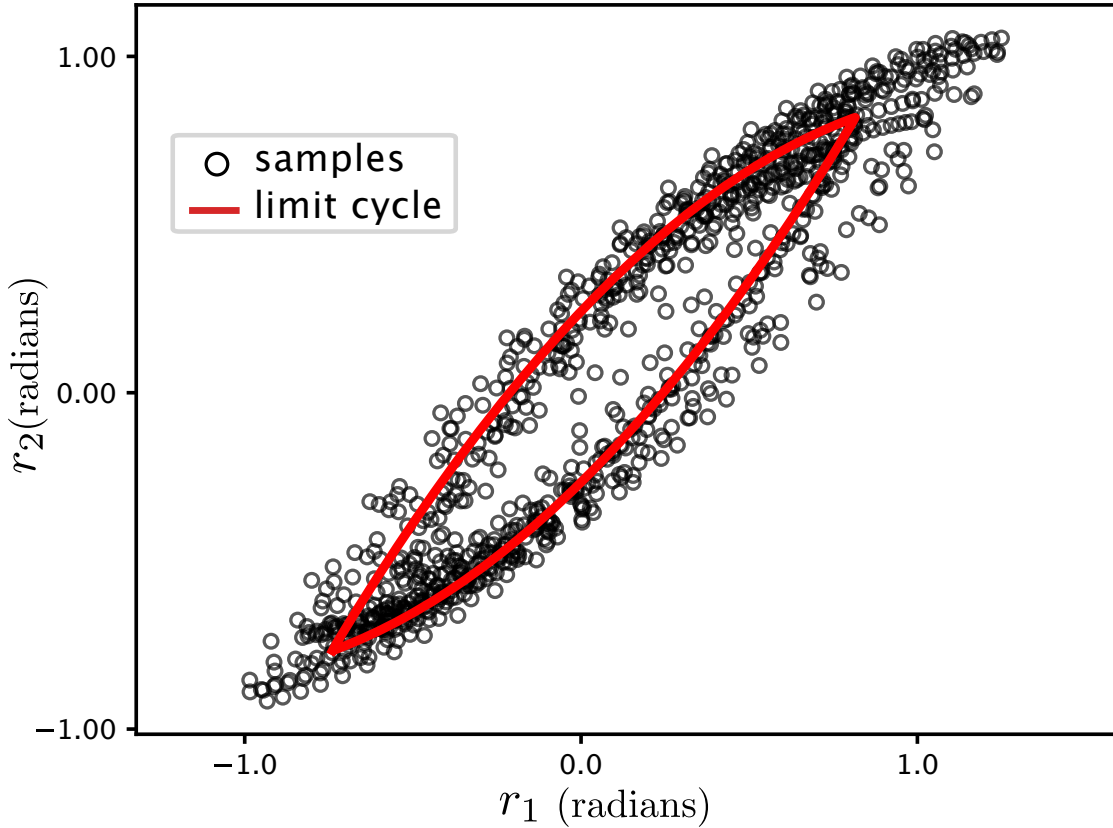


Parameter	Range
$T_{\text{low}}$	20 ~ 41 °C
$T_{\text{high}}$	45 ~ 65 °C
$t_{\text{cool}}$	2 ~ 8 hrs
$t_{\text{heat}}$	0.5 ~ 3 hrs
$\eta_{\text{cool}}$	$\frac{1}{32} \sim 1$
$\eta_{\text{heat}}$	$\frac{1}{32} \sim 1$

**Table 4-II.** Full input parameter ranges.

### 4.5.5 Shape Parameterization

The compliant nature of the devices and external forces makes the internal shape high-dimensional. However, fitting models to a highly dimensional shape space will likely cause overfitting. We thus seek a reduced-order representation of the shape of the system. Here, principal components analysis (PCA) is a simple candidate reduction method that could serve this purpose. We tried PCA on the streamline along the crawler body, from which we calculated the internal angles between each of the segments. Fitting a PCA model on the internal angles, we found that the first two principal components (modes) reconstruct over 90% of the data. While it is a straightforward way of reducing the effective degrees of freedom, the complex coupling between segments led to principal components that lacked a clear physical interpretation. As an alternative, we considered that the volume of each active section is a more physical, descriptive candidate for representing the system shape variables. To do this, we estimated the volume at each time point based on contour points from the FEA simulation, and it was used to calculate the enclosed area (2D volume for our planar FEA). This parameterization provided a clear, intuitive relationship between the two segments, and exhibited the phase lag between the smaller and larger bilayer segments that we expected from the design.



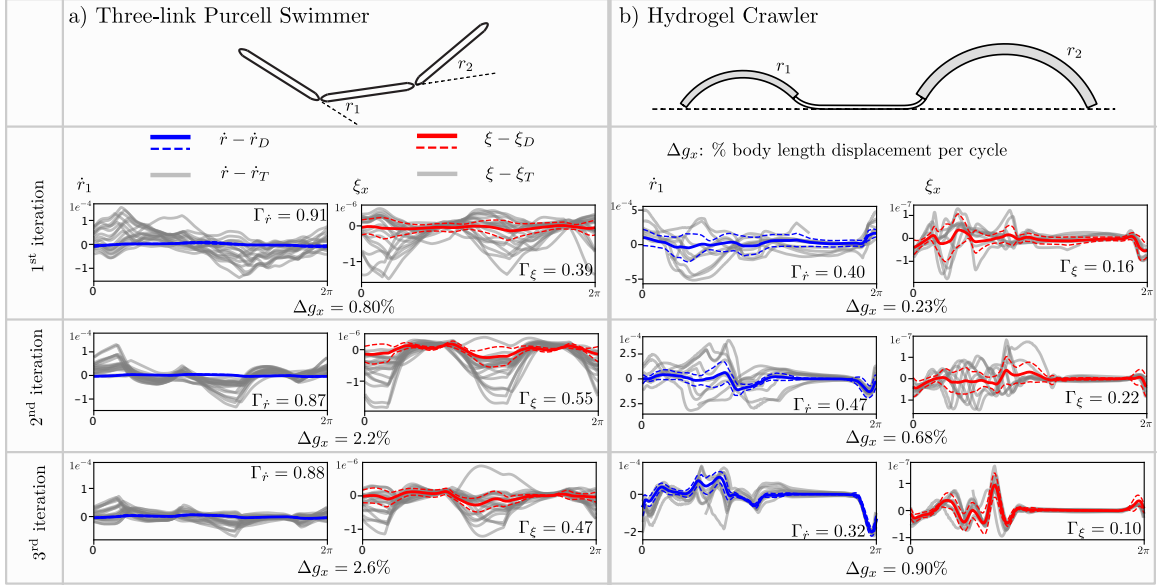
**Figure 4-5.** Sampled shape data points around a limit cycle. Note that we cannot directly sample the full shape and shape velocity space due to the presence of the actuator dynamics. Fortunately, adding perturbation to the input parameters does generate shape and shape velocity variations around the limit cycle, which is essential for building models of behaviors.

#### 4.5.6 Input Optimization

As a demonstration, we optimized the input parameters to maximize the displacement per cycle. The objective function is defined as

$$F(u) = \Delta g_x - \lambda t_{\text{cycle}}, \quad (4.10)$$

where  $\Delta g_x$  is the displacement in the  $x$  direction per cycle, and  $t_{\text{cycle}}$  is the cycle time.  $\lambda$  is a penalty factor that controls the trade-off between the above two terms. The objective function is to maximize the net displacement per cycle. During the optimization process, we noticed that the optimizer tended to find cycles with the



**Figure 4-6.** Model prediction metrics for the Purcell swimmer (top left) and hydrogel crawler (top right) model. We plot sample model prediction errors in phase coordinates from each of the model refining iterations. Also, the average displacement, in terms of percentage body length per actuation cycle, each iteration is shown below the plots. The metrics are calculated on the testing set which is unseen by the model both on shape velocity predictions and body velocity predictions. Test set trajectories (grey) along with the mean (solid) and standard deviation (dashed) of the data-driven model prediction error are plotted for each iteration. The body velocity prediction error is inherently propagated from the shape velocity prediction error, and thus the body velocity prediction errors are unsurprisingly greater than the shape velocity prediction error. While the model prediction errors start to drop as we shrink the sampling range due to overfitting, the average behavioral objective improves each iteration.

longest possible cycle times, optimizing cycle-to-cycle distance, with no penalty on time, pushing the results toward the boundary. To address this, we added a regularizing term to penalize the cycle time.

We started by sampling 100 points (resulting in 100 cycles of system motion) in the full input parameter space as shown in Table 4-II. We performed ten-fold cross-validation to avoid overfitting. A rough model of the system was built using the initially sampled data points, and then the model was used to optimize for an input parameter that maximizes the objective function above. The optimization was performed using the Sequential Least Squares Programming algorithm where

the local gradient and Hessian were estimated using finite differencing. Once the numerical optimum was obtained, we shrank the range of each input parameter by 35%, centering at the optimum, and repeated the optimization process. We repeated this process three times, and the parameters converged to a performant gait. The model improvement metrics were calculated for each iteration as shown in Fig. 4-6

## 4.6 Discussion and Conclusion

In this work, we designed and implemented a data-driven modeling framework for dissipative systems with low-bandwidth actuator dynamics. We showed the success of this method in predicting behaviors on a classical simplified model, the Purcell swimmer, with a modified class of passive shape dynamics. We built on prior work, relaxing the requirement that at least one shape element is accessed through high-bandwidth control. In doing so, our method enables the modeling of novel mechanisms like the hydrogel crawler, whose internal degrees of freedom all exhibit a passive response to controllable stimuli. Despite the challenges often associated with designing control signals under low-bandwidth actuation, the robot was intentionally designed to capitalize on morphology-induced actuation asymmetries for locomotion. We showed not only that we could model the crawler with accuracy beyond the phase-averaged gait, but that the system was capable of using this model in a gradient-based optimization scheme to rapidly identify a viable crawling maneuver. The broader implications of this result are that we now have a rational framework to pursue data-driven modeling and optimization of a much larger class of underactuated systems. For applications in biology, where continuous, soft interfaces facilitate safe interaction with the body, this method provides the potential to model new mechanisms pre-deployment in the body and even *in situ*, since variation amongst morphology and environment across patients can be significant. Key additional future efforts will include power, actuation, and sensing at the scales desired for the locomotion application. In addition, it is

well known that there is a significant phase lag between muscle activation and body movements [140], suggesting that our approach can be used to better understand the underlying neural control problem [141].

The model improvement saturated and decayed as the sampling region was reduced in the hydrogel crawler gait optimization result. This likely occurred because there was less variation in the sampled data (our tube of gait distribution data had a smaller overall volume). This is loosely analogous to convergence results in adaptive control, which often rely on sufficient excitation of the dynamics; in a similar vein, we do not expect to learn informative improvement without cycles that excite significant dynamic variation.

We have shown that the model was iteratively improved and the optimizer settles at a performant gait. Given computational power, sampling from a variety of initial conditions could excite a variety of achieved locally optimal gaits, from which a more globally optimal gait could be selected. In this project, a single cycle of FEA simulation takes hours to run, making brute forcing multiple initial conditions infeasible. Rather than characterizing this as a limitation, we argue that our approach captures its ideal utilization. Many next-generation soft robots will be hard to model a priori. Imagine custom morphologies deployed in uncharacterized patient bodies (such as an artificial heart valve [142]). This mechanism must leverage the ability to build a viable control policy from very little data. Likewise, the objectives for the hydrogel crawler were simply to obtain a functional control policy from very little data. We have provided a framework through which it and similar robots can rapidly obtain functional motion primitives. We juxtapose our method with reinforcement learning [143], deep learning architectures [144], and gait optimization [145] which require large amounts of data that are unavailable in simulation and impractical to empirically obtain for many applications (e.g. in biomedical settings).

The viable gaits achieved through this sample efficient optimization could practi-

cally extend to the real world. While many samples might be available in simulated environments, there are many platforms that must be system-identified in the field. *In-situ* system identification (such as the type we implement here) paired with a gradient-based optimizer provides a tractable, system-oriented way to pursue optimization of robot behaviors in hard-to-model environments. Further efforts can be made to wrap control, modeling and optimization into an in-the-loop process, where the robot can learn to navigate in complex, time-varying environments without human intervention.

# Chapter 5

## Discussion and Future Work

In this dissertation, we have proposed solutions to some practical challenges in implementing data-driven geometric methods. We have shown that the proposed methods are effective in solving the problems in a variety of scenarios, including real-time system identification, enabling steering via gait control, and optimizing control signals for a hydrogel crawler.

Chapter 2 takes its inspiration from data-driven modeling, geometric mechanics, and adaptive control techniques. We implement an adaptive system identification framework and demonstrate its effectiveness in improving the performance of primarily kinematic locomotion systems, such as those dominated by dissipative forces or zero momentum conservation. We highlight the adaptive model’s ability to seamlessly adapt to various terrains and modify behaviors iteratively within a behavior optimization framework. This capability not only enhances fundamental behaviors but also enables precise motion tracking. Additionally, we showcase its successful application in refining behaviors, recovering from injury, and adapting to different terrains. The methods are shown to excel in scenarios where simulations provide inadequate guidance for real-world situations.

This dissertation has mainly focused on the development of data-driven geometric methods for analytical systems such as Purcell swimmers. To further validate these

methods, we plan to apply them to real systems, with a hexapod robot (RHex) as a potential candidate. We have done preliminary work on this topic in simulation environments, and real-world experiments are needed to further validate the proposed methods.

In Chapter 3, we introduce both local and global gait morphing algorithms designed to modify a default gait for single-parameter steering control. Using a simplified swimmer as an example, we perform a numerical comparison of these two approaches, revealing that the local method, while suboptimal, remains effective for modest steering control. An advantage of the local approach is its adaptability to soft robots or other systems where local approximations of constraint curvature can be derived from data, as creating an exact global model may be infeasible. This local approach makes the steering control methods applicable to the adaptive geometric modeling framework presented in Chapter 2, where such data-driven models are only local in nature.

Chapter 4 presents a novel method that constructs a series of interconnected models, encompassing actuator and locomotor dynamics derived from data collected during stochastically perturbed and repeated behaviors. These connected models are based on general formulations of dissipative Lagrangian systems with symmetry, making them applicable to a wide range of robots with first-order, low-pass actuator dynamics, including hydrogel crawlers driven by swelling-based actuators. These models accurately capture the dynamics of system shape and body movements in a simplified swimming robot model. Furthermore, we apply this approach to a stimulus-responsive hydrogel simulator, which emulates the intricate chemo-mechanical interactions responsible for shape changes in biomedically relevant micromachines. We propose a method for numerically optimizing control signals through iterative refinement of models, which we apply to optimize the input waveform for the hydrogel crawler. This transfer to real-world environments holds promise for applications in locomotor design and biomedical engineering.



Many systems such as the hydrogel crawler have complex internal dynamics that are difficult to model due to a number of factors, including imprecise manufacturing, complex geometry, and unknown composite material properties. In such cases, first-order, low-pass actuator dynamics are less likely to capture the system’s true dynamics. A more dedicated approach such as methods presented in [146] could potentially be useful in capturing the internal dynamics of the system to further improve the model capabilities. In addition, the tools introduced in this dissertation are limited to principally kinematic systems; future works could utilize increasingly powerful machine learning techniques to capture momentum effects as well.

The framework of geometric mechanics is becoming more and more inclusive in terms of its applicable scenarios, including some dynamical systems that were previously deemed to fall outside of its scope. A major advancement of this dissertation is that it provides useful tools for the implementation of data-driven geometric methods that make them particularly suitable for real-world situations. Obtaining steering control is only one of many potential applications of the adaptive data-driven geometric methods. We see the potential of integrating our tools into more general artificial intelligence algorithms to facilitate robots to autonomously learn and adapt in complex scenarios. In particular, the real-time capability of efficiently capturing system behaviors can be useful in popular control algorithms, such as nonlinear model predictive control [147] and adaptive control techniques [148], to help robots make real-time decisions autonomously. For example, perception can provide certain expectations in changes in the robot-environment interactions, which can be used to tune the internal learning modules (adaptive models) to be either more reactive or robust. On a larger scale, we anticipate the potential of our tools to enable stronger connections among subfields of robotics like perception, planning, modeling, and control, bringing us one step closer to achieving animal-like locomotion in robots.

# References

1. Schneier, M., Schneier, M. & Bostelman, R. Literature review of mobile robots for manufacturing (2015).
2. Katevas, N. *Mobile robotics in healthcare* (IOS Press, 2001).
3. Davids, A. Urban search and rescue robots: from tragedy to technology. *IEEE Intelligent systems* **17**, 81–83 (2002).
4. Fujita, M. & Kitano, H. Development of an autonomous quadruped robot for robot entertainment. *Autonomous robots* **5**, 7–18 (1998).
5. Hardarson, F. *Locomotion for difficult terrain* (Citeseer, 1998).
6. Koditschek, D. E. & Bühler, M. Analysis of a Simplified Hopping Robot. *The International Journal of Robotics Research* **10**, 587–605 (1991).
7. Park, J. H. & Kim, K. D. *Biped robot walking using gravity-compensated inverted pendulum mode and computed torque control* in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No. 98CH36146)* **4** (1998), 3528–3533.
8. Raibert, M. H. *Legged robots that balance* (MIT press, 1986).
9. Saranli, U., Buehler, M. & Koditschek, D. E. RHex: A simple and highly mobile hexapod robot. *The International Journal of Robotics Research* **20**, 616–631 (2001).
10. Hirose, S. & Fukushima, E. F. Snakes and strings: New robotic components for rescue operations. *The International Journal of Robotics Research* **23**, 341–349 (2004).
11. Anderson, J. M., Streitlien, K., Barrett, D. & Triantafyllou, M. S. Oscillating foils of high propulsive efficiency. *Journal of Fluid mechanics* **360**, 41–72 (1998).
12. Donley, J. M., Sepulveda, C. A., Konstantinidis, P., Gemballa, S. & Shadwick, R. E. Convergent evolution in mechanical design of lamnid sharks and tunas. *Nature* **429**, 61–65 (2004).
13. Ma, K. Y., Chirarattananon, P., Fuller, S. B. & Wood, R. J. Controlled flight of a biologically inspired, insect-scale robot. *Science* **340**, 603–607 (2013).
14. Shepherd, R. F., Ilievski, F., Choi, W., Morin, S. A., Stokes, A. A., Mazzeo, A. D., Chen, X., Wang, M. & Whitesides, G. M. Multigait soft robot. *Proceedings of the National Academy of Sciences* **108**, 20400–20403 (2011).
15. Tangorra, J., Phelan, C., Esposito, C. & Lauder, G. Use of Biorobotic Models of Highly Deformable Fins for Studying the Mechanics and Control of Fin Forces in Fishes. *Integrative and Comparative Biology* **51**, 176–189 (June 2011).

16. Tesch, M., Schneider, J. & Choset, H. *Using response surfaces and expected improvement to optimize snake robot gait parameters in 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2011), 1069–1074.
17. Galloway, K. C., Haynes, G. C., Ilhan, B. D., Johnson, A. M., Knopf, R., Lynch, G. A., Plotnick, B. N., White, M. & Koditschek, D. E. X-RHex: A highly mobile hexapedal robot for sensorimotor tasks (2010).
18. Hyun, D. J., Seok, S., Lee, J. & Kim, S. High speed trot-running: Implementation of a hierarchical controller using proprioceptive impedance control on the MIT Cheetah. *The International Journal of Robotics Research* **33**, 1417–1445 (2014).
19. Cheng, M.-Y. & Lin, C.-S. *Genetic algorithm for control design of biped locomotion in 1995 IEEE International Conference on Systems, Man and Cybernetics. Intelligent Systems for the 21st Century* **2** (1995), 1315–1320.
20. Furusho, J. & Masubuchi, M. Control of a dynamical biped locomotion system for steady walking (1986).
21. Goswami, A. Postural stability of biped robots and the foot-rotation indicator (FRI) point. *The International Journal of Robotics Research* **18**, 523–533 (1999).
22. McGeer, T. *Passive walking with knees in Proceedings., IEEE International Conference on Robotics and Automation* (1990), 1640–1645.
23. Collins, S. H., Wisse, M. & Ruina, A. A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research* **20**, 607–615 (2001).
24. Collins, S., Ruina, A., Tedrake, R. & Wisse, M. Efficient bipedal robots based on passive-dynamic walkers. *Science* **307**, 1082–1085 (2005).
25. Hürmüzlü, Y. & Moskowitz, G. D. The role of impact in the stability of bipedal locomotion. *Dynamics and Stability of Systems* **1**, 217–234 (1986).
26. Westervelt, E. R., Grizzle, J. W. & Koditschek, D. E. Hybrid zero dynamics of planar biped walkers. *IEEE transactions on automatic control* **48**, 42–56 (2003).
27. Westervelt, E. R., Grizzle, J. W., Chevallereau, C., Choi, J. H. & Morris, B. *Feedback control of dynamic bipedal robot locomotion* (CRC press, 2018).
28. Raibert, M., Blankespoor, K., Nelson, G. & Playter, R. Bigdog, the rough-terrain quadruped robot. *IFAC Proceedings Volumes* **41**, 10822–10825 (2008).
29. Seipel, J. & Holmes, P. Three-dimensional translational dynamics and stability of multi-legged runners. *The International Journal of Robotics Research* **25**, 889–902 (2006).
30. Duan, X., Chen, W., Yu, S. & Liu, J. *Tripod gaits planning and kinematics analysis of a hexapod robot in 2009 IEEE International Conference on Control and Automation* (2009), 1850–1855.
31. Roy, S. S. & Pratihar, D. K. Kinematics, dynamics and power consumption analyses for turning motion of a six-legged robot. *Journal of Intelligent & Robotic Systems* **74**, 663–688 (2014).
32. Jayne, B. C. What defines different modes of snake locomotion? *Integrative and comparative biology* **60**, 156–170 (2020).

33. Hirose, S. *Biologically Inspired Robots: Snake-Like Locomotors and Manipulators* in (1993).
34. Burdick, J. W., Radford, J. & Chirikjian, G. S. A 'sidewinding' locomotion gait for hyper-redundant robots. *Advanced Robotics* **9**, 195–216 (1994).
35. Miller, G. *Snake robots for search and rescue, neurotechnology for biomimetic robots* in *Proc. Int. Conf. Robotics and Automation (ICRA)* (2002), 16–21.
36. Yamada, H. *S. development of amphibious snake-like robot acm-r5* in *The 36th international symposium on robotics (ISR 2005), Tokyo* (2005).
37. Wright, C., Johnson, A., Peck, A., McCord, Z., Naaktgeboren, A., Gianfortoni, P., Gonzalez-Rivero, M., Hatton, R. & Choset, H. *Design of a modular snake robot* in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2007), 2609–2614.
38. Gravish, N. & Lauder, G. V. Robotics-inspired biology. *Journal of Experimental Biology* **221**, jeb138438 (2018).
39. Aguilar, J., Zhang, T., Qian, F., Kingsbury, M., McInroe, B., Mazouchova, N., Li, C., Maladen, R., Gong, C., Travers, M., *et al.* A review on locomotion robophysics: the study of movement at the intersection of robotics, soft matter and dynamical systems. *Reports on Progress in Physics* **79**, 110001 (2016).
40. Kovač, M. The bioinspiration design paradigm: A perspective for soft robotics. *Soft Robotics* **1**, 28–37 (2014).
41. Ijspeert, A. J. Biorobotics: Using robots to emulate and investigate agile locomotion. *science* **346**, 196–203 (2014).
42. Kelly, S. D. & Murray, R. M. Geometric phases and robotic locomotion. *Journal of Robotic Systems* **12**, 417–431 (1995).
43. Marsden, J. E. & Ostrowski, J. Symmetries in motion: Geometric foundations of motion control. *National Academies Press* (1998).
44. Yezzi, A. J. & Soatto, S. Deformation: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision* **53**, 153–167 (2003).
45. Shapere, A. & Wilczek, F. *Geometric phases in physics* (World scientific, 1989).
46. Marsden, J. E. & Ratiu, T. S. *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems* (Springer Science & Business Media, 2013).
47. Ostrowski, J. P., Desai, J. P. & Kumar, V. Optimal Gait Selection for Nonholonomic Locomotion Systems. *The International Journal of Robotics Research* **19**, 225–237 (2000).
48. Bloch, A. M. *Nonholonomic Mechanics and Control* (Springer New York, NY, 2003).
49. Simple mechanical control systems with constraints. *IEEE Transactions on Automatic Control* **45**, 1420–1436 (2000).
50. Ostrowski, J. P., Desai, J. P. & Kumar, V. Optimal gait selection for nonholonomic locomotion systems. *The International journal of robotics research* **19**, 225–237 (2000).

51. Bullo, F. & Lynch, K. M. Kinematic controllability for decoupled trajectory planning in underactuated mechanical systems. *IEEE Transactions on Robotics and Automation* **17**, 402–412 (2001).
52. Hatton, R. L. & Choset, H. *Connection vector fields for underactuated systems in 2008 2nd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics* (2008), 451–456.
53. Kelly, S. D., Pujari, P. & Xiong, H. *Geometric mechanics, dynamics, and control of fishlike swimming in a planar ideal fluid in Natural Locomotion in Fluids and on Surfaces: Swimming, Flying, and Sliding* (2012), 101–116.
54. Melli, J. B., Rowley, C. W. & Rufat, D. S. Motion planning for an articulated body in a perfect planar fluid. *SIAM Journal on applied dynamical systems* **5**, 650–669 (2006).
55. Avron, J. E. & Raz, O. A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin. *New Journal of Physics* **10**, 063016 (2008).
56. Shammas, E. A., Choset, H. & Rizzi, A. A. Natural gait generation techniques for principally kinematic mechanical systems (2005).
57. Hatton, R. L. & Choset, H. Geometric motion planning: The local connection, Stokes’ theorem, and the importance of coordinate choice. *The International Journal of Robotics Research* **30**, 988–1014 (2011).
58. Hatton, R. L. & Choset, H. Nonconservativity and noncommutativity in locomotion. *The European Physical Journal Special Topics* **224**, 3141–3174 (2015).
59. Hatton, R. L. & Choset, H. Geometric Swimming at Low and High Reynolds Numbers. *IEEE Transactions on Robotics* **29**, 615–624 (2013).
60. Dai, J., Faraji, H., Gong, C., Hatton, R. L., Goldman, D. I. & Choset, H. *Geometric Swimming on a Granular Surface in Robotics: Science and Systems* (2016).
61. Ramasamy, S. & Hatton, R. L. *Soap-bubble optimization of gaits in 2016 IEEE 55th Conference on Decision and Control (CDC)* (2016), 1056–1062.
62. Choi, J., Bass, C. & Hatton, R. L. *Optimal Gait Families using Lagrange Multiplier Method in 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2022), 8873–8878.
63. Krose, B. *An introduction to neural networks* (1996).
64. Quinlan, J. R. Induction of decision trees. *Machine learning* **1**, 81–106 (1986).
65. Kaelbling, L. P., Littman, M. L. & Moore, A. W. Reinforcement learning: A survey. *Journal of artificial intelligence research* **4**, 237–285 (1996).
66. Sjöberg, J., Zhang, Q., Ljung, L., Benveniste, A., Delyon, B., Glorennec, P.-Y., Hjalmarsson, H. & Juditsky, A. Nonlinear black-box modeling in system identification: a unified overview. *Automatica* **31**, 1691–1724 (1995).
67. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W. & Abbeel, P. *Domain randomization for transferring deep neural networks from simulation to the real world in 2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (2017), 23–30.
68. Bittner, B., Hatton, R. L. & Revzen, S. Geometrically optimal gaits: a data-driven approach. *Nonlinear Dynamics* **94**, 1933–1948 (2018).

69. Revzen, S. *Neuromechanical control architectures of arthropod locomotion* (University of California, Berkeley, 2009).
70. Kvalheim, M. D., Bittner, B. & Revzen, S. Gait modeling and optimization for the perturbed Stokes regime. *Nonlinear Dynamics* **97**, 2249–2270 (2019).
71. Bittner, B., Hatton, R. L. & Revzen, S. Data-driven geometric system identification for shape-underactuated dissipative systems. *Bioinspiration & Biomimetics* **17**, 026004 (2022).
72. Deng, S., Liu, J., Datta, B., Pantula, A., Gracias, D. H., Nguyen, T. D., Bittner, B. A. & Cowan, N. J. A Data-Driven Approach to Geometric Modeling of Systems with Low-Bandwidth Actuator Dynamics. *arXiv preprint arXiv:2307.01062* (2023).
73. Zhao, D., Bittner, B., Clifton, G., Gravish, N. & Revzen, S. Walking is like slithering: A unifying, data-driven view of locomotion. *Proceedings of the National Academy of Sciences* **119**, e2113222119 (2022).
74. Chong, B., He, J., Li, S., Erickson, E., Diaz, K., Wang, T., Soto, D. & Goldman, D. I. Self-propulsion via slipping: Frictional swimming in multilegged locomotors. *Proceedings of the National Academy of Sciences* **120**, e2213698120 (2023).
75. Pantula, A., Datta, B., Shi, Y., Wang, M., Liu, J., Deng, S., Cowan, N. J., Nguyen, T. D. & Gracias, D. H. Untethered unidirectionally crawling gels driven by asymmetry in contact forces. *Science Robotics* **7**, eadd2903 (2022).
76. McClung, A. J. *Techniques for dynamic maneuvering of hexapedal legged robots* (ProQuest, 2006).
77. Task-level control of rapid wall following in the American cockroach. *Journal of Experimental Biology* **209**, 1617–1629 (2006).
78. Seipel, J. E., Holmes, P. J. & Full, R. J. Dynamics and stability of insect locomotion: a hexapedal model for horizontal plane motions. *Biological cybernetics* **91**, 76–90 (2004).
79. Proctor, J. & Holmes, P. Steering by transient destabilization in piecewise-holonomic models of legged locomotion. *Regular and Chaotic Dynamics* **13**, 267–282 (2008).
80. Cowan, N. J., Ma, E. J., Cutkosky, M. & Full, R. J. in *Robotics Research* 540–550 (Springer, Siena, Italy, 2003).
81. Zhao, D. & Revzen, S. Multi-legged steering and slipping with low DoF hexapod robots. *Bioinspiration & biomimetics* **15**, 045001 (2020).
82. Deng, S., Hatton, R. L. & Cowan, N. J. *Enhancing Maneuverability via Gait Design* in *Proc IEEE Int Conf Robot Autom* (2022), 5799–5805.
83. Deng, S., Cowan, N. J. & Bittner, B. A. Adaptive Gait Modeling and Optimization for Principally Kinematic Systems. *arXiv preprint arXiv:2310.02141* (2023).
84. Shi, R., Chen, K.-L., Fern, J., Deng, S., Liu, Y., Scalise, D., Huang, Q., Cowan, N. J., Gracias, D. H. & Schulman, R. Shape-shifting microgel automata controlled by DNA sequence instructions. *bioRxiv*, 2022–09 (2022).
85. Grandia, R., Taylor, A. J., Ames, A. D. & Hutter, M. *Multi-layered safety for legged robots via control barrier functions and model predictive control* in *2021 IEEE International Conference on Robotics and Automation (ICRA)* (2021), 8352–8358.

86. Gong, Y. & Grizzle, J. W. Zero dynamics, pendulum models, and angular momentum in feedback control of bipedal locomotion. *Journal of Dynamic Systems, Measurement, and Control* **144**, 121006 (2022).
87. Weingarten, J. D., Lopes, G. A., Buehler, M., Groff, R. E. & Koditschek, D. E. Automated gait adaptation for legged robots in *IEEE International Conference on Robotics and Automation, 2004* **3** (2004), 2153–2158.
88. Brambilla, G., Buchli, J. & Ijspeert, A. J. Adaptive four legged locomotion control based on nonlinear dynamical systems in *International Conference on Simulation of Adaptive Behavior* (2006), 138–149.
89. Sun, C., He, W., Ge, W. & Chang, C. Adaptive neural network control of biped robots. *IEEE transactions on systems, man, and cybernetics: systems* **47**, 315–326 (2016).
90. Peng, X. B., Berseth, G., Yin, K. & Van De Panne, M. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG)* **36**, 1–13 (2017).
91. Heess, N., Tb, D., Sriram, S., Lemmon, J., Merel, J., Wayne, G., Tassa, Y., Erez, T., Wang, Z., Eslami, S., *et al.* Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286* (2017).
92. Tan, J., Zhang, T., Coumans, E., Iscen, A., Bai, Y., Hafner, D., Bohez, S. & Vanhoucke, V. Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332* (2018).
93. Mullins, G. E., Stankiewicz, P. G., Hawthorne, R. C. & Gupta, S. K. Adaptive generation of challenging scenarios for testing and evaluation of autonomous vehicles. *Journal of Systems and Software* **137**, 197–215 (2018).
94. Rambabu, R., Vadakkepat, P., Tan, K. C. & Jiang, M. A mixture-of-experts prediction framework for evolutionary dynamic multiobjective optimization. *IEEE transactions on cybernetics* **50**, 5099–5112 (2019).
95. Tsuda, B., Tye, K. M., Siegelmann, H. T. & Sejnowski, T. J. A modeling framework for adaptive lifelong learning with transfer and savings through gating in the prefrontal cortex. *Proceedings of the National Academy of Sciences* **117**, 29872–29882 (2020).
96. Rozen-Levy, S., Messner, W. & Trimmer, B. A. The design and development of branch bot: a branch-crawling, caterpillar-inspired, soft robot. *The International Journal of Robotics Research* **40**, 24–36 (2021).
97. Pantula, A., Datta, B., Shi, Y., Wang, M., Liu, J., Deng, S., Cowan, N. J., Nguyen, T. D. & Gracias, D. H. Untethered unidirectionally crawling gels driven by asymmetry in contact forces. *Science Robotics* **7**, eadd2903 (2022).
98. Li, M., Xi, N., Wang, Y. & Liu, L. Progress in nanorobotics for advancing biomedicine. *IEEE Transactions on Biomedical Engineering* **68**, 130–147 (2020).
99. Park, C., Ozturk, C. & Roche, E. T. Computational design of a soft robotic myocardium for biomimetic motion and function. *Advanced Functional Materials* **32**, 2206734 (2022).
100. Webster-Wood, V. A., Guix, M., Xu, N. W., Behkam, B., Sato, H., Sarkar, D., Sanchez, S., Shimizu, M. & Parker, K. K. Biohybrid robots: Recent progress, challenges, and perspectives. *Bioinspiration & Biomimetics* (2022).

101. Cendra, H., Marsden, J. E. & Ratiu, T. S. Geometric mechanics, Lagrangian reduction, and nonholonomic systems. *Mathematics Unlimited—2001 and Beyond*, 221–273 (2001).
102. Kelly, S. D. & Murray, R. M. *The geometry and control of dissipative systems in Proceedings of 35th IEEE Conference on Decision and Control* **1** (1996), 981–986.
103. Hayes, M. H. *Statistical digital signal processing and modeling* (John Wiley & Sons, 1996).
104. Purcell, E. M. Life at low Reynolds number. *American journal of physics* **45**, 3–11 (1977).
105. Cox, R. G. The motion of long slender bodies in a viscous fluid Part 1. General theory. *Journal of Fluid Mechanics* **44**, 791–810 (1970).
106. Williams, G., Aldrich, A. & Theodorou, E. A. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics* **40**, 344–357 (2017).
107. Carron, A., Arcari, E., Wermelinger, M., Hewing, L., Hutter, M. & Zeilinger, M. N. Data-driven model predictive control for trajectory tracking with a robotic arm. *IEEE Robotics and Automation Letters* **4**, 3758–3765 (2019).
108. Hogan, F. R. & Rodriguez, A. Reactive planar non-prehensile manipulation with hybrid model predictive control. *The International Journal of Robotics Research* **39**, 755–773 (2020).
109. Ye, C., Ma, S., Li, B. & Wang, Y. *Turning and side motion of snake-like robot in IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004* **5** (2004), 5075–5080.
110. Lee, J., Sponberg, S. N., Loh, O. Y., Lamperski, A. G., Full, R. J. & Cowan, N. J. Templates and anchors for antenna-based wall following in cockroaches and robots. *IEEE Trans Robot* **24**, 130–143 (2008).
111. Sfakiotakis, M. & Tsakiris, D. Biomimetic Centering for Undulatory Robots. *International Journal of Robotics Research* (2007).
112. Murray, R. & Sastry, S. S. *Nonholonomic Motion Planning: Steering Using Sinusoids* tech. rep. UCB/ERL M91/45 (EECS Department, University of California, Berkeley, 1991).
113. Bittner, B. & Revzen, S. Optimizing Gait Libraries via a Coverage Metric. *arXiv:2107.08775*. eprint: [arXiv:2107.08775](https://arxiv.org/abs/2107.08775) (2021).
114. Hatton, R. L., Burton, L. J., Hosoi, A. E. & Choset, H. *Geometric maneuverability with applications to low reynolds number swimming in 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2011), 3893–3898.
115. Ramasamy, S. & Hatton, R. L. The Geometry of Optimal Gaits for Drag-Dominated Kinematic Systems. *IEEE Transactions on Robotics* **35**, 1014–1033 (2019).
116. Becker, L. & Stone, H. On self-propulsion of micro-machines at low Reynolds number: Purcell’s three-link swimmer. *Journal of Fluid Mechanics* **490**, 15–35 (2003).
117. Full, R. J. & Koditschek, D. E. Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *Journal of Experimental Biology* **202**, 3325–3332 (1999).



118. Hatton, R. & Choset, H. *Approximating displacement with the body velocity integral* in (2009).
119. Hatton, R. L. & Choset, H. in *Robotics: Science and Systems VII* 113–120 (2012).
120. Hatton, R. L., Dear, T. & Choset, H. Kinematic Cartography and the Efficiency of Viscous Swimming. *IEEE Transactions on Robotics* **33**, 523–535 (2017).
121. Flanders, H. Differentiation Under the Integral Sign. *American Mathematical Monthly* **80**, 615–627 (1973).
122. Liu, J., Erol, O., Pantula, A., Liu, W., Jiang, Z., Kobayashi, K., Chatterjee, D., Hibino, N., Romer, L. H., Kang, S. H., Nguyen, T. D. & Gracias, D. H. Dual-Gel 4D Printing of Bioinspired Tubes. *ACS Applied Materials & Interfaces* **11**, 8492–8498 (2019).
123. Ramasamy, S. & Hatton, R. L. Optimal gaits for drag-dominated swimmers with passive elastic joints. *Physical Review E* **103**, 032605 (3 2021).
124. Trivedi, D., Rahn, C. D., Kier, W. M. & Walker, I. D. Soft robotics: Biological inspiration, state of the art, and future research. *Applied Bionics and Biomechanics* **5**, 99–117 (2008).
125. Li, S., Bai, H., Liu, Z., Zhang, X., Huang, C., Wiesner, L. W., Silberstein, M. & Shepherd, R. F. Digital light processing of liquid crystal elastomers for self-sensing artificial muscles. *Science Advances* **7**, eabg3677 (2021).
126. TolleyMichael, T., ShepherdRobert, F., GallowayKevin, C., WoodRobert, J., WhitesidesGeorge, M., *et al.* A resilient, untethered soft robot. *Soft Robotics* (2014).
127. Liu, B., Ozkan-Aydin, Y., Goldman, D. I. & Hammond, F. L. *Kirigami skin improves soft earthworm robot anchoring and locomotion under cohesive soil* in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)* (2019), 828–833.
128. Aydin, Y. O., Molnar, J. L., Goldman, D. I. & Hammond, F. L. *Design of a soft robophysical earthworm model* in *2018 IEEE International Conference on Soft Robotics (RoboSoft)* (2018), 83–87.
129. Zhao, H., Li, Y., Elsamadisi, A. & Shepherd, R. Scalable manufacturing of high force wearable soft actuators. *Extreme Mechanics Letters* **3**, 89–104 (2015).
130. Lu, H., Zhang, M., Yang, Y., Huang, Q., Fukuda, T., Wang, Z. & Shen, Y. A bioinspired multilegged soft millirobot that functions in both dry and wet conditions. *Nature Communications* **9**, 3944 (2018).
131. Chen, W.-H., Yang, W., Peach, L., Koditschek, D. E. & Sung, C. R. Kinigami: Algorithmic Design of Compliant Kinematic Chains From Tubular Origami. *IEEE Transactions on Robotics* **39**, 1260–1280 (2 Oct. 11, 2022). published.
132. Ozkan Aydin, Y., Chong, B., Gong, C., Rieser, J. M., Rankin, J. W., Michel, K., Nicieza, A. G., Hutchinson, J., Choset, H. & Goldman, D. I. *Geometric mechanics applied to tetrapod locomotion on granular media* in *Biomimetic and Biohybrid Systems: 6th International Conference, Living Machines 2017, Stanford, CA, USA, July 26–28, 2017, Proceedings 6* (2017), 595–603.
133. Breger, J. C., Yoon, C., Xiao, R., Kwag, H. R., Wang, M. O., Fisher, J. P., Nguyen, T. D. & Gracias, D. H. Self-folding thermo-magnetically responsive soft microgrippers. *ACS applied materials & interfaces* **7**, 3398–3405 (2015).

134. Han, D., Farino, C., Yang, C., Scott, T., Browe, D., Choi, W., Freeman, J. W. & Lee, H. Soft robotic manipulation and locomotion with a 3D printed electroactive hydrogel. *ACS applied materials & interfaces* **10**, 17512–17518 (2018).
135. Na, H., Kang, Y.-W., Park, C. S., Jung, S., Kim, H.-Y. & Sun, J.-Y. Hydrogel-based strong and fast actuators by electroosmotic turgor pressure. *Science* **376**, 301–307 (2022).
136. Revzen, S. & Guckenheimer, J. M. Estimating the phase of synchronized oscillators. *Physical Review E* **78**, 051907 (5 2008).
137. Revzen, S. & Kvalheim, M. *Data driven models of legged locomotion in Micro- and Nanotechnology Sensors, Systems, and Applications VII* (eds George, T., Dutta, A. K. & Islam, M. S.) **9467** (SPIE, 2015), 94671V.
138. *ABAQUS/Standard, Version 2021* English. United States, 2021.
139. Chester, S. A. & Anand, L. A thermo-mechanically coupled theory for fluid permeation in elastomeric materials: Application to thermally responsive gels. *Journal of the Mechanics and Physics of Solids* **59**, 1978–2006 (2011).
140. Ding, Y., Sharpe, S. S., Wiesenfeld, K. & Goldman, D. I. Emergence of the advancing neuromechanical phase in a resistive force dominated medium. *Proceedings of the National Academy of Sciences* **110**, 10123–10128 (2013).
141. Madhav, M. S. & Cowan, N. J. The synergy between neuroscience and control theory: the nervous system as inspiration for hard control challenges. *Annual Review of Control, Robotics, and Autonomous Systems* **3**, 243–267 (2020).
142. Hasan, A., Saliba, J., Pezeshgi Modarres, H., Bakhaty, A., Nasajpour, A., Mofrad, M. R. & Sanati-Nezhad, A. Micro and nanotechnologies in heart valve tissue engineering. *Biomaterials* **103**, 278–292 (2016).
143. Erden, M. S. & Leblebicioğlu, K. Free gait generation with reinforcement learning for a six-legged robot. *Robotics and Autonomous Systems* **56**, 199–212 (2008).
144. Tsounis, V., Alge, M., Lee, J., Farshidian, F. & Hutter, M. DeepGait: Planning and Control of Quadrupedal Gaits Using Deep Reinforcement Learning. *IEEE Robotics and Automation Letters* **5**, 3699–3706 (2020).
145. Winkler, A. W., Bellicoso, C. D., Hutter, M. & Buchli, J. Gait and Trajectory Optimization for Legged Systems Through Phase-Based End-Effector Parameterization. *IEEE Robotics and Automation Letters* **3**, 1560–1567 (2018).
146. Hidir, E. K., Uyanik, I. & Morgül, Ö. Harmonic transfer functions based controllers for linear time-periodic systems. *Transactions of the Institute of Measurement and Control* **41**, 2171–2184 (2019).
147. Grüne, L., Pannek, J., Grüne, L. & Pannek, J. *Nonlinear model predictive control* (Springer, 2017).
148. Åström, K. J. & Wittenmark, B. *Adaptive control* (Courier Corporation, 2013).