Data-Driven Methods for Geometric Systems

by

Brian A. Bittner

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Robotics)
in The University of Michigan
2021

Doctoral Committee:

Associate Professor Shai Revzen, Chair
Professor Anthony M. Bloch
Professor Jessy W. Grizzle
Associate Professor Ross L. Hatton, Oregon State University

Brian A. Bittner

babitt@umich.edu

ORCID iD: 0000-0001-8694-4317

To my family

ACKNOWLEDGEMENTS

I have many people to thank for my time in Ann Arbor. Foremost, I am grateful for the thoughtful mentorship of my advisor Shai Revzen. He helped me build on my foundations as an engineer to understand the philosophy and value of scientific research. His capability to be hands-on and hands-off at the right time allowed me to learn and explore in a way that made my studies very engaging. I admire his breadth of knowledge across research disciplines and have appreciated his contagious curiosity and passion for problem solving. I was fortunate to enter the lab when there was a collection of strong senior graduate students. Matt Kvalheim and George Council went out of their way to take me under their wing in the BIRDS Lab. It was a pleasure to collaborate with Matt on a component of my thesis, which could not have been accomplished without his depth of mathematical insights. George was my neighbor in the lab for four years, and we regularly discussed research perspectives and ideas. These conversations were often spontaneous but could go on for hours. I missed these conversations during my last year of school. To my labmates Dan Zhao, Marion Anderson, Ziyou Wu, Vikram Sachdeva, and Taylor McLaughlin, it has been a pleasure to work and learn with you. I am thankful for collaborations with Ross Hatton, Suresh Ramasamy, Glenna Clifton, Nick Gravish, and Eleni Gourgou. I enjoyed learning from your perspectives and expertise. Thanks as well to David Remy and Ram Vasudevan for insightful feedback on this project during my qualifications exam. At conferences, I remember engaging in discussions with students Benjamin Shih, Julian Whitman, Yves Nazon, and Kevin Green.

I began at the Robotics Institute when it was beginning its third year as an institution. It was special to feel as though I was contributing, in my own small way, to the foundation of a research and learning community that will be a leader in robotics education for years to come. From beginning to end, I felt strongly supported by the institute's leadership, namely Jessy Grizzle and Ella Atkins. Special thanks to Denise Edmonds for her support as my academic advisor. As a younger student becoming involved in the robotics community through student government and other means, it was nice to have strong mentorship at the student level. Mia Stevens, Katie Skinner, Ross Hartley, Alexandra Carlson, Kevin Lieberman, and Nils Smit-Anseeuw made me feel included immediately and helped me find my niche within the student community. I can only hope to have paid forward to others what their mentorship meant to me. Thanks as well to students Grant Gibson and Eva Mungai for organizing Michigan's student-run Legged Locomotion Seminar. It was one of my favorite things to participate in. I became friends with Paarth Shah, Yu Ming Chen, and Brandt Monson through acting and control classes. I appreciated their positivity and sense of humor and am grateful for their friendship.

Pittsburgh is where my passion for robotics started. Thanks to my father, a testing and validation engineer in field robotics, for taking me to hang out with him at work when I was twelve. Thanks to Nicolas Vandapel and Dan Munoz for allowing me to explore coding and robotics in a research and development context during high school. I'd like to thank Koushil Sreenath, Steve Collins, George Kantor, Howie Choset, and Matt Travers for being strong sources of faculty mentorship in robotics during my undergraduate years. Thanks to Trevor Decker for encouraging me to participate in research and make friends in the robotics student community. Thanks to Quan Nguyen for giving me the confidence that I could succeed as a graduate student. Thanks to Arun Srivatsan for taking me under his wing at my first conference.

I am very grateful for the support of my family and friends. It was always really nice to feel like people wanted to know how my research projects were going when I was visiting. I have enjoyed being able to share that part of my life with you all. To my parents, I am thankful for your sacrifice to open doors for me and your unconditional support and encouragement. I love you and am grateful to have you in my life. To Sean, Jake, Dan, and Aidan, your candor and sense of humor make life colorful. To my partner Julia, I love you and am lucky to have your endless support. When I met you during my first summer in Ann Arbor, I had no idea what experiences we would come to share and cherish.

Finally, thanks to my committee and the scholars that came before me. I have valued the opportunity to peek over the shoulders of giants.

TABLE OF CONTENTS

CHAPTER

LIST OF FIGURES

Figure

ABSTRACT

The tools of geometric mechanics provide a compact representation of locomotion dynamics as “the reconstruction equation”. We have found this equation yields a convenient form for estimating models directly from observation data. This convenience draws from the method’s relatively rare feature of providing high accuracy models with little effort. By little effort, we point to the modeling process’s low data requirements and the property that nothing about the implementation changes when substituting robot kinematics, material properties, or environmental conditions, as long as some intuitive baseline features of the dynamics are shared. We have applied data-driven geometric mechanics models toward optimizing robot behaviors both physical and simulated, exploring robots’ ability to recover from injury, and efficiently creating libraries of maneuvers to be used as building blocks for higher-level robot tasks. Our methods employed the tools of data-driven Floquet analysis, providing a phase that we used as a means of grouping related measurements, allowing us to estimate a reconstruction equation model as a function of phase in the neighborhood of an observed behavior. This tool allowed us to build models at unanticipated scales of complexity and speed. Our use of a perturbation expansion for the geometric terms led to an improved estimation procedure for highly damped systems containing non-trivial but non-dominating amounts of momentum. Analysis of the role of passivity in dissipative systems led to another extension of the estimation procedure to robots with high degrees of underactuation in their internal shape, such as soft robots. This thesis will cover these findings and results, simulated and physical, and the surprising practicality of data-driven geometric mechanics.

CHAPTER I

Introduction

1.1 Motivation

The field of robotics has succeeded in producing a variety of capable platforms for exploration. These range from wheeled vehicles to platforms that walk, crawl, slither, swim, and fly. In many cases, the roboticist adds a degree of freedom to the mechanism, expanding its range of expressions. The added joint's potential to create new, better, safer, and more useful behaviors is inhibited by the added exponential increase in sampling complexity that is now required to catalog and cache the expanded space of robot maneuvers. Here the infamous *curse of dimensionality* inhibits the roboticist from designing an arbitrarily complex robot and immediately arriving at a portfolio of useful behaviors for locomotion. Algorithms that can handle behavior discovery and optimization on complex robots have long been of interest to robotics researchers, and challenges in their development persist today.

Some methods, including blackbox nonliear optimization [Wei+02; Cal+14] and reinforcement learning [Gu+17; Haa+18; Ha+20] considered the robot to have no exploitable dynamical structure. During the optimization of behaviors, an action of the robot is conventionally represented by some parametrization. These methods treated the sampling of a behavior or gait[1] parametrization as a direct experiment on the

---

[1] A gait is a rhythmic behavior.

robot. Here, the purpose of each experiment is to measure some performance metric that is being optimized. Some typical metrics include speed, efficiency, accuracy, and stability. These methods possess the advantage of getting real-world experimental validation with each experiment. However, these methods are sampling inefficient because they never use a dynamical model to inform the search for useful behaviors. In the time it took to sample one gait on the robot, thousands of offline samples could have been passed through a dynamical model to simulate (predict) what the result of the experiment would have been. Researchers have used neural nets to model dynamical systems from data, providing the capability to sample trajectories offline [Ban+16; Nag+18]. The neural net data-driven approximation of the dynamics could then guide the search for a useful policy or control strategy on the robot. As more parameters and layers are added to a given neural net, its capacity to represent more general dynamical systems behavior improves. The cost of this generality feature is the tendency to overfit low quality models to datasets of limited size. Unfortunately, there is very limited structural knowledge in the field of machine learning to inform how much shrinking of a neural net is required to reduce overfitting, or how large a neural net must be to maintain its ability to capture complex system behaviors.

Other robotics researchers attempted to use dynamics and physical models as a structural guide for the optimization of behaviors. In the extreme, roboticists modeled the robot via submission of the full body configuration to a finely tuned physics engine [Tod11; Hee+17; Ana+18; Mae+18]. These model intensive methods allowed, as mentioned above, for the sampling of many trajectories offline. For model intensive behavior optimization algorithms, the solutions produced are fundamentally limited by the fidelity of the model used. When these models were accurate with respect to real-world robot encounters, the methods were quite successful. However, when true features of the robot interaction were not captured by the model, it substantially limited the utility of these methods. The challenge of extending simulation based results to the real world is well documented [Sün+18; Tan+18; MJD18].

Deriving models from first principles (such as by Newtonian, Lagrangian, or Hamiltonian mechanics) allowed for the use of well developed methods for control and stability analysis [BLM; Wes+18] and even the analytical derivation of optimal gaits [HC11a]. Some gait optimization methods yielded successful results on robots such as bipeds using the full complexity of first principles models [Kaj+; GG17]. Typically, methods using analytical models require precise accounting of the shape and mass distribution of the robot components. As complexity increases, analysis of the model becomes more opaque and specific to the system. If this complexity involves adding a degree of freedom, behavioral optimization becomes more sampling intensive.

A theory for reduced-order control architectures in robots and animals [FK99; Sei+17] provided a way to think about useful model reductions for these systems. Analogous to the case of dropping layers and parameters in neural nets, reduced-order models here involve dropping degrees of freedom in the mechanical model or introducing coarse representations of the robot's mass distribution or morphology. Some popular simplistic models (termed "templates") for legged systems include the lateral leg spring (LLS) model [SH00b; SH00a], and the spring loaded inverted pendulum (SLIP) model [BF93]. Here, reductions allow for a lower space of control inputs and behavioral outputs to be studied, making analysis and sampling more tractable. For example, researchers analytically explored the passive dynamics of reduced-order models of bipedal and quadrupedal gaits to generate a continuum of stable gaits spanning various walking strategies [Gan+15; Gan+18]. However, just as in the case of neural nets, it is formally unclear when these reduced-order, first principles models lose the representational capacity critical to predict accurate motions on the physical hardware. Some current research explores optimizing the ability of a reduced-order model [CP20] or pair of reduced or models [GP07] to generate useful modeling and control architectures for higher-order systems. The use of Gaussian processes [DR11] and representation learning frameworks [BZL06] suffer a similar issue. It is often unclear what level of reduction is acceptable to fully represent even a local neighborhood of the dynamics.

In this thesis work, we provided a framework for data-driven modeling where reductions are informed by geometry and physics. The first reduction involved the well understood Lagrangian reduction under symmetry [CMR01; Blo+96]. This formulation used homogeneity in the environment to remove dependence on the position of the robot from the dynamics. The second reduction involved the insight that as the role of generalized momentum in the dynamics is minimized, the dynamics collapse to a first-order model captured by a *mechanical connection* [MO98]. Through quotienting position dependence and second-order dynamics from the space of motion model constraints, we maintained precise knowledge of what representational capacity is lost. Employing the tools of data-driven Floquet analysis, we obtained a phase-sample pairing that we used as a means of grouping related measurements. This allowed us to estimate a reconstruction equation model as a function of phase in the neighborhood of an observed behavior. By using this approach, we fit a data-driven model to system behavior that is both **general** to the expressive capacity of the systems being analyzed and **compact** enough to quickly fit observation data without overfitting. The ability to system identify quickly and the clarity of the required physical assumptions involved in the reduction make this an attractive method for roboticists using field-deployed systems. It is a low effort, high accuracy system identifier. The cost of access to use these rare qualities is to be able to specify the relative roles of inertia and damping in the system, as well as to specify the existence of environmental symmetries. In order to use the tools developed in this work, the robot must have an interaction with the environment that is not dependent on where it is placed in the environment. Additionally, friction must dominate inertia in the dynamics of the system. We will specify more concretely the mathematical meaning of these statements, but they remain physically intuitive.

In this thesis, we approach the *curse of dimensionality* by building data-driven models of behaviors in their operating environment that are highly sample efficient (we show results where models are built on eight degree of freedom systems with 30 cycles of data). We took advantage of the ability to compute behavioral models with small amounts of data to inform iterative behavioral modifications via a policy gradient approach. This method allowed robots to optimize various navigation-related goal functions within minutes, making such strategies realistic for use in the field. Next, we will discuss the background of modeling systems through the perspectives of geometric mechanics and oscillator theory. The technical tools developed there provide the structure for the key contributions to data-driven modeling presented in this thesis.

1.2 Background

1.2.1 Geometric Mechanics

In the field of geometric mechanics, a system's dynamical constraints and group symmetries can generate a reduced system [CMR01; Blo+96]. The dynamical constraints are typically derived from Lagrangian or Hamiltonian conservation laws, whereas group symmetries typically extend from symmetry in the environment. These reduced systems and their consequently simplified models allow deeper insight and analysis to the behavior of systems (such as the snakeboard [Ost96]) whose equations take an otherwise opaque form, requiring inspection by numerical simulation [Ost+]. Writing control systems with the structure of a reduced Lagrangian [OB98], the equations of motion can be split into separate pieces consistent with the geometric features induced by the symmetries of the system. This controllable form has value to the field of robotic locomotion, where the equations for internal variables, position, and momentum are all factored as dependent but distinct equations.

Analysis of the equations leads to a simple observation that zero momentum systems can encounter a subsequent reduction. Systems whose pfaffian constraints fully determine a locomotion map can be further reduced to a first-order dynamical system, a *mechanical connection* [MO98] that linearly maps internal shape velocities to body velocities of the system. This class of systems includes those with no initial momentum operating under momentum conservation laws (such as a floating astronaut [MO98] or high Reynolds swimmer [HC13]). For these examples, starting from rest in the absence of a gravitational field eliminates the ability to start or accumulate momentum. Another way to arrive at a connection is to add a Rayleigh dissipation function that dominates the dynamics. In the case of the swimmer, adding dominating drag forces can push it into the low Reynolds regime [HC13].

The work of [HC11a; HC15; RH19] used the insights provided by the structure of the *mechanical connection* to develop new approaches to coordinate system selection, gait identification, and behavioral optimization. The work in those papers used a global analytical model. However, for real-world animals and robots, high fidelity global models are often challenging to obtain. System identification techniques [HC13; Dai+; Sch+19] allowed for data-driven model extraction via a large sampling of animal or robot maneuvers. A limitation of these methods is that in the field, one can often not afford to spend resources collecting observation data. Modeling techniques that are lean on data requirements can be used and deployed more easily. Furthermore, requiring little data to update or re-compute models enables a modeling approach to be more adaptive to environmental and internal changes. One way to reduce data requirements of previous work may involve modeling the connection in the neighborhood of a behavior, rather than trying to measure more global properties of the connection. In the next section, we discuss an approach to modeling rhythmic behaviors from observation data.

1.2.2 Data-Driven Oscillator Theory

Animal and robot locomotion often exhibit rhythmic qualities. Like in [Sei+17], we consider a rhythmic behavior to be one that has an exponentially stable periodic solution in the configuration space. A phase oscillator offers a very simplistic model for this type of behavior, where a phase points to a value on a cycle in the configuration space (including the internal shape and velocity in the body frame). The projection of a complex oscillating system to the reduced-order structure of a cycle inherited some perspective from [FK99], which outlined a theory of reduced-order control architectures in animals and robots. This level of simplicity is valuable for the analysis of many systems, but more complex models can capture more detailed features of system locomotion. A more general asymptotically stable oscillator for the system can be hypothesized, such that a differential equation governs the dynamics around a gait (cycle on the phase oscillator) within some stability basin [Sei+17]. For systems that meet these criteria, some natural questions to ask when observing locomotion data are: What is the period of the gait? What is the average behavior of an oscillation with respect to phase? What are the dynamics of attraction to this average behavior? Data-Driven Floquet Analysis provides closed form solutions to these answers [Rev09]. The estimator assumes a general, asymptotically stable oscillator, which has been shown to apply even for some classes of hybrid systems [BRS15; RK15a]. While this generality in representation can lead to models that can be applied to a broad class of systems, large quantities of data are typically needed to compute predictive models. If a system were to exhibit only first-order dynamical structure, such as the *mechanical connections* of the previous section, it would likely reduce the data requirements needed to compute these predictive oscillator models.

Another contribution from data-driven oscillator theory is the development of algorithms for the estimation of phase. There are heuristic examples of phase assignment which are applied to various use cases, such as prosthetic limbs [Gre+14]. A data-driven phase estimator called phaser [RG08] computes the asymptotic phase map given raw oscillation data and specification of a Poincare section. This asymptotic phase map estimates what phase any data point would have if it were projected onto its value on a phase oscillator. For the most recent developments on improvements to this estimation, see [Kva09].

One use of phaser is for the grouping of noisy data to build phase-ordered models of robot actions. It can also be used to investigate hypotheses about the neuro-mechanical control architectures of animals [RKF09]. The inspiration of these tools included various hypotheses about the role of phase feedback in locomotion, such as phase reduction in the nervous system [HC85] and the effects of the structure of coupled oscillators on central pattern generators [Gol+99].

1.3 Contributions

1.3.1 Chapter 2: Designed a sample efficient estimator for the connection governing a periodic motion

The technical contributions of this chapter are the construction of the data-driven geometric modeling framework and the subsequent data-driven geometric gait optimizer. We showed in this chapter that for systems with the structure of a *mechanical connection* (reduced lagrangian with no momentum), noisy robot behaviors can be modeled from data in a highly sample efficient way. We built these models by applying data-driven oscillator theory and assuming the first-order dynamical structure of the mechanical connection. This resulted in a data-driven geometric modeling technique that can be plugged into a broad class of dissipative systems acting in homogenous environments.

1.3.2 Chapter 2: Designed a policy gradient method that uses the connection estimator

We complemented the ability to compute local models in the vicinity of behaviors with the development of a policy gradient optimizer. We call this framework data-driven geometric gait optimization, detailing useful ways to parametrize principally kinematic gaits and step through the parameter space in between optimization iterations.

1.3.3 Chapter 2: Validated the optimization method on simulated Purcell swimmers and physical wheeled snake robots

Using the three-link and nine-link Purcell swimmer, we showed that with 30 cycles of following an extremal gait under nontrivial noise perturbations, high fidelity models could be computed from the data for the swimmers. This validated the capacity of the modeling framework to yield predictive models for systems of small and high dimension. The data-driven geometric gait optimizer was validated through its consistent convergence to high efficiency behaviors, provided a wide variety of initial gaits and levels of perturbation influence. The validation of these tools is detailed in the online publication [BHR18]. Using the same algorithm, we showed that the implementations of [BHR18] were successful on hardware platforms in §2.9, such as the five-link and nine-link wheeled snake and various other robots. The modeling and optimization of simulated systems covered in this Chapter were published in Springer Nonlinear Dynamics [BHR18].

1.3.4 Chapter 3: Developed a primitive library optimization metric termed coverage

Individual performance criteria, such as robustness and efficiency, often drive behavior selection in mobile robotic platforms. However, navigational capabilities can rely heavily on the coordination of all behaviors available to the robot, rather than the performance of one behavior. Here, we studied how a collection of available gaits (cyclic internal motions) can be related to the variability and density of external motions that a system can achieve on its position space represented as a Lie group. We attempted to capture this capability, termed "coverage", with an associative fast-to-compute cost function that can be used to simultaneously optimize a collection of gaits. The cost function can be weighted such that users may specify the importance of various classes of external motions. The technical contribution of this chapter is the development of the coverage metric for gait optimization. We highlighted the advantage of valuing a collection of behaviors for the qualities of their composition rather than their individual qualities.

1.3.5 Chapter 3: Validated this optimization on Purcell swimmers that experience injury and a robot made of tree branches

We designed two unconventional viscous swimming robots, called a two-slider swimmer and the three-branch swimmer. The two-slider swimmer cannot translate without rotating while executing a gait that does not intersect itself. The three-branch swimmer lacks the bilateral symmetry typically seen in robotic platforms. We showed that both of these systems are able to achieve a collection of useful poses for navigation by composing a small collection of gaits. This demonstrated the value of querying the ability of a gait library to cover a space of local pose variations rather than the individual qualities of the gaits. By placing this new metric into the data-driven geometric gait optimization framework, we presented the ability of the Purcell swimmer to find high coverage collections of gaits under a variety of circumstances. We tested convergence to high coverage gait libraries while changing the number of links of the swimmer and the initial gaits available to the optimizer [BR19]. We extended these results by taking the swimmers' post coverage optimization and locking the joint exhibiting maximal amplitude. We saw that swimmers with four or more joints are able to quickly recover a high coverage collection of gaits of similar quality. Finally, we plugged a robot made of tree branches into the framework. The system was able to find a collection of gaits for navigation with just 10 minutes and 24 seconds of experimental data [BR20a].

1.3.6 Chapter 4.1: Designed a sample efficient estimator for shape-underactuated dissipative systems

The technical contribution here is the extension of the data-driven geometric modeling tool from fully actuated dissipative systems to shape-underactuated systems. We propose this for modeling of systems like soft robots, which typically have a handful of control inputs and a high dimensional internal state. We also showed that the control inputs for these systems could be converted to the space of forces and torques rather than shape velocities. This is more compatible with the control of actuators that are typically used in soft systems.

1.3.7 Chapter 4.1: Validated the model accuracy of the estimator on a collection of simulated shape-underactuated dissipative systems

We showed that a first-order Taylor expansion of the geometric model, with light assumptions on the passive dynamics, contains a different collection of regression terms than a naive first-order and second-order Taylor expansion of the dynamics. The proposed approximation contains more precise regressors than the naive first-order Taylor expansion. The proposed approximation also contains regressors that scale linearly with the number of passive elements, rather than the quadratic growth seen in the naive second-order approximation. The ability of the proposed approximation to remain compact as the number of passive elements increases could have strong ramifications for soft robotic systems, which typically have a high degree of underactuation. We demonstrated the ability of the proposed approximation to model a variety of shape-underactuated viscous swimming platforms. We started with the simple swimmers that can only achieve linear motion and built our way up to systems with six underactuated degrees of freedom (via the nine-link Purcell swimmer). Chapter 4.1 is avaiable on arXiv [BHR20].

1.3.8 Chapter 4.2: Designed a sample efficient estimator for systems in the perturbed Stokes regime

Our collaborator Kvalheim applied results from singular perturbation theory to show that when momentum decays quickly, it exponentially converges to a function of shape and shape velocity. Thus, momentum is an output rather than a state variable. Our contribution to the work was to extend the class of regressors available for modeling the connection so that it could capture this additional term.

1.3.9 Chapter 4.2: Validated that the perturbed Stokes models are beneficial in intermediate Reynolds numbers

Using a simulation of a paddleboat over Reynolds numbers ranging from $\frac{1}{10}$ to 10 we demonstrated that the perturbed Stokes estimation offers consistently higher predictive quality than the connection estimation. We showed that this beneficial performance is consistent for nominal gaits and gaits that are extremal in the viscous limit. We also showed that the benefits of the perturbed Stokes model persisted when adding joints to the system. The results of Chapter 4.2 were published in Springer Nonlinear Dynamics[KBR19].

1.4 Discussion

In this thesis, we demonstrate the ability to rapidly engineer behaviors for a variety of complex systems. These have numerous important ramifications for the field of robotics. Because we model the robots over a space of dynamical constraints relevant to dissipative systems, robots can be made of many shapes, sizes, materials, and actuation methodologies and still be modeled by the exact same algorithm. This altogether avoids the cumbersome burden of supplying the morphology, mass distribution, and compliance properties to a simulator or first principles modeling framework. This yields low effort, high accuracy modeling for the engineer that could facilitate new levels of performance for field robots, such as in the following examples. Robotic fleets can be deployed in the field for exploration tasks and persist autonomously through substantial environmental changes (gravel to sand) or internal injuries (motor failures). The ability to self-model (Chapters II and IV) and produce a compact set of behaviors for useful navigation policies (Chapter III) can enable systems to quickly get back up and running. These methods naturally extend to the use of imprecisely fabricated robots acting in locations with uncharacterized environmental properties. Robots made by some nontraditional fabrication methods can include ones made of paper and tape [Fit+17], soft robots made of fastened pneumatic actuators [Bru+20], and dielectric elastomers [SBC15]. For these systems, the ability to system identify quickly allows for the construction and deployment of robots without precision bookkeeping of their mechanical parameters. Furthermore, robots are getting smaller, often for the purposes of biomedical applications [Ric+17]. These systems are expected to experience low Reynolds interaction with the environment due to their size. Other biomedical applications, such as an artificial heart muscle, appear to act in friction dominated environments [Par+20]. These robots have complicated sources of actuation, and their operating environment can vary greatly. Having efficient sampling algorithms that make behavior engineering procedural and time-efficient could open new possibilities in the field of biomedical engineering.

The methods of this thesis are suitable for systems that are imprecise in both fabrication and control. A strength of the algorithms presented is their ability to model high noise behaviors. These methods are informed by experiences in biology [Rev09], where organisms are known to experience substantial variation in internal motion while executing cyclic behaviors like walking, crawling, and slithering. Methods for the assignment of phase [RG08] permit organization and analysis of noisy behaviors. The use of such tools positions our methods to be used to model both animal behaviors with low precision or robotic systems with low gain control. By building intelligent algorithms that admit design and control sub-optimalities, we suggest an expansion of the robotics community's interpretation of what constitutes a practically useful mechanism. Cost effective robots are more likely to make their capabilities accessible to a broader audience, mitigating financial barriers to engaging with valuable technologies, whether these values lie in healthcare, education, manufacturing, transportation, or leisure.

A major advantage of the methods of this thesis is that it makes robotic engineering with low cost hardware more accessible. While the methods result from rigorous dynamical systems analysis, the assumptions for implementation are physically intuitive in nature. An understanding of friction, inertia, environmental homogeneity, and basic trajectory design are the primary requirements for successfully engineering systems by this methodology. Understanding of advanced physics, control theory, inverse kinematics, Denavit-Hartenberg tables, hard and soft constraints, and hyperparameter tuning are not necessary to use these tools.

Lastly, we will touch upon how the tools of this thesis may allow us to test new hypotheses in biomechanics and evolution. The results from Chapter IV §4.2 give us a way to think about how features of the dynamics change as the roles of inertia and friction are modified. This allows us to ask interesting questions about the evolution of locomotion strategies and morphologies in biological organisms. Earth's earliest organisms may have been microswimmers, inhabiting friction dominated regimes. If so, as branches of their evolutionary trees began to leverage inertia as a contributor to locomotion, what changes occurred in the system morphology and nervous system? Were there some changes that were particularly advantageous from the geometric viewpoint? These are questions we now have the tools to test hypotheses about. For some of today's biological movers, we observe thought provoking trends. Agile animals like gazelles appear to have less mass distribution in their appendages, removing some complexity and nuance that could otherwise result in the dynamics on p. For example, we can see through principles of momentum conservation that, during an aerial phase, low mass appendages can diminish the impact of shape change on change in the angular momentum of the center of mass. This is desirable from the perspective of building an estimator for body velocity that has a smaller number of dependent variables. Does this make it easier to learn behavioral strategies for perching and leaping? For example, is there a consequently simpler organization of the neural system that relates to feedback for maneuver execution? The tools of this thesis are limited to the Stokes and perturbed Stokes regime, but future work could extend these estimators to systems like birds and gazelles, where momentum plays a more significant role in important behaviors. Using data-driven geometric mechanics, we can build the ability to assess the degree of complication in the dynamics in a way that can be compared across systems of vastly different morphologies. Robotics offers a testbed to further explore such hypotheses and ascertain which geometric principles can be embedded usefully in mechanism design. Moving to the friction dominated regime, animals like snakes and octopi appear to have maneuvers that are dominated by contact and friction forces rather than inertia. The ramifications of this on the dynamical model, highlighted by the geometric formulation, are that the simple, Jacobian-like *mechanical connection* model governs the dynamics. By being friction dominated, the system avoids any interaction with the complexity-inducing momentum. These systems happen to have much more widely distributed mass than their agile counterparts. According to data-driven geometric mechanics, there is no body velocity estimator complexity cost to mass distribution in friction governed systems. With data-driven geometric mechanics, we have a toolbox for systems engineering, but also a toolbox to investigate hypotheses about the advantages of various biological behaviors and morphologies.

CHAPTER II

Data-Driven Modeling and Optimization in the Stokes Regime

2.1 Motivation

The ability to move effectively through the environment is both a defining property of animals and a highly desirable capability for man-made systems such as robots and vehicles. Locomotion (aquatic, terrestrial, and aerial) is most commonly achieved by having a moving body change shape in a way that produces reaction forces from the environment; these reaction forces in turn propel the body. A key question in both robotics and animal research is thus: Does a given gait cycle optimally exploit this propulsive relationship, and if not, what changes to the gait would improve its performance?

This paper details a new approach to answering these questions, by presenting a practical extension of geometric gait optimization theory that incorporates techniques from the data-driven modeling of gaits as oscillators. By efficiently producing a local geometric mechanics model of the observed motion, we can then employ this model to rapidly evaluate the gradient of a goal function with respect to gait parameters. Because this performance simulation is very fast, the number of gait parameters being optimized can be so large that estimating such a gradient by direct experimentation is nigh impossible; in this chapter we give an example with 264 parameters optimized in 30 trials of 30 cycles each.

The framework presented in this paper is made possible by combining work in two fields that have developed largely in parallel. In the field of geometric mechanics, Hatton has developed a framework for characterizing gait efficiency in terms of the length and area of the cycle in the shape space [HC11b; HC15; HDC17; RH16; RH17]. Applying these principles to systems that lack an analytical model remains an open area of investigation, especially when high-dimensionality makes the exhaustive exploration of system dynamics from [Hat+13; Dai+] infeasible, or when considering an animal whose motions we cannot directly command.

In the field of oscillator theory, Revzen developed a set of tools for extracting oscillator-like motion models from noisy and irregularly-spaced data [Rev09; RG08; RG11]. In addition to the method's robustness to the intrinsic system noise of biological and physical systems, it extends well to high dimensional shape spaces. A limitation, however, is the lack of insight that these models provide for gait improvements.

Applying the data-driven oscillator and geometric approaches together enhances their respective capabilities: the data-driven oscillator tools can provide the geometric models with the specific information needed for evaluating a performance criterion and its gradient, improving their predictive power relative to the quantity of available data. Conversely, by viewing the system as a mechanical connection (as a opposed to a general second-order dynamical system) the data-driven oscillator models can ignore certain aspects of the system dynamics that are irrelevant to the optimality of the gait, thus significantly reduce the algorithmic complexity of model extraction.

Here, we lay out a framework for combining the geometric insights from Hatton's work with the data-driven oscillator model construction from Revzen's work. Our combined approach uses noise in the dynamics of a system that follows a nominal gait cycle to build a model of the system dynamics in the neighborhood of this gait cycle. Inserting this model into our geometric tools then provides estimates of both how optimal this gait is relative to nearby cycles and what perturbations can be applied to the cycle to best improve its performance. This estimation technique has two primary use cases.

- The first is as a tool that allows for verification of postulated goal functions for observed animal locomotion.
- The second is in field robotics, where an efficient, noise resistant gait optimization algorithm can potentially enable learning effective gaits without requiring precise analytical models of the robot or its interactions with the environment.

The approach presented in this paper offers a collection of advantages in speed, scalability, and model reduction for the estimation of motion models and subsequent optimization of gaits. These advantages derive from the use of geometric mechanics models governed by a connection (the "principal kinematic case" in the language of [OB98]). The absence of a momentum term in the equation of motion implies that the contribution of different segments of motion do not strongly depend on each other, allowing motion models to be integrated in parallel instead of sequentially in time. With multi-processor computing becoming cheaper, this offers the opportunity for dramatic speedups in the computation of motion plans. Additionally, connections expose the fact that the systems they govern are, for practical purposes, half the dimension of general mechanical systems. As the very name “geometric mechanics” suggests, in these systems the geometry of motion in body shape space governs the outcome of motions, admitting a description with only one dimension per degree of freedom, instead of the two needed in conventional Newtonian mechanics. Despite this great promise of geometric mechanics models, little work has been done on producing them in a data-driven way. Developing a tool for the data-driven creation of such

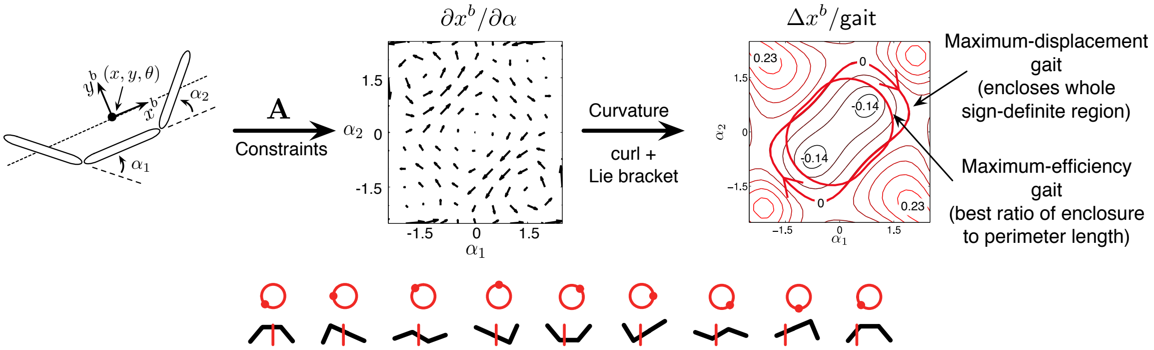

Figure 2.1: Key elements of the geometric paradigm, drawn from to [HC11b; HDC17]. Given a locomoting system (left), the system dynamics and constraints produce a relationship between changes in shape and changes in position (center). When the system executes a cyclic change in shape (bottom), the net displacement induced by this gait corresponds to how much curvature of the constraints the gait encompasses, and the time-effort cost of executing this gait is the length of the path it traces out in the shape space (right). In the bottom animation, the top row indicates the phase of the swimmer in a gait, with the shape at each phase shown in the bottom row. The motion can be viewed from left-to-right or right-to-left, with the vertical bar (red) serving as a static reference point.

models allows us to explore their value for both scientific and engineering applications.

Below we review the geometric and data-driven approaches, and then synthesize them into a tool for simultaneously estimating and optimizing locomotion models. Using simulated mechanical swimming platforms, we illustrate the precision of these data-driven geometric mechanics models, and demonstrate that optimal gaits can be learned with very few trials. Finally, we discuss the utility of the new methods for both system identification and field robotics.

2.2 Geometry of Locomotion

The first thread of prior work that this paper draws upon is geometric modeling of locomotion. When analyzing a mobile deformable system, it is convenient to separate its configuration space Q (i.e. the space of its generalized coordinates q) into a position space G and a shape space R, such that the position $g \in G$ locates the system in the world, and the shape $r \in R$ gives the relative arrangement of the particles that compose it.[1]

During locomotion, changes in the system's shape provoke reaction forces from the environment that in turn drive changes in the system's position. For the purposes of this paper, we adopt a (geometric) locomotion model

$$\overset{\circ}{g} = \mathbf{A}(r)\dot{r}, \tag{2.1}$$

where $\mathbf{A}$, the *local connection*, linearly maps the shape velocity $\dot{r}$ to the body velocity $\overset{\circ}{g} = g^{-1}\dot{g}$ (i.e., the position velocity in the body frame's current forward, lateral, and rotational directions). The local connection acts similarly to the Jacobian map of a kinematic mechanism — it takes the velocity of joints to the position velocity (here, of the body frame instead of end effector) that they generate under the constraints imposed on the system.

We model the cost of changing shape as corresponding to the length s of the trajectory through the shape space,

$$s = \int \sqrt{dr^{\mathsf{T}} \mathcal{M}(r)\, dr} = \int_0^T \sqrt{\dot{r}^{\mathsf{T}} \mathcal{M}(r)\, \dot{r}}\; dt, \tag{2.2}$$

where $\mathcal{M}$ is a Riemannian metric on the shape space that weights the costs of changing shape in different directions.

This connection-and-metric model applies to systems that move by pushing directly against their environment with negligible accumulated momentum in "gliding" modes, and whose energetic costs are dominated by internal or external dissipative effects. This model has been analytically derived for swimmers in viscous fluids [AR08; HDC17], and experimentally validated for several robots in dry granu-

---

[1] In the parlance of geometric mechanics, this assigns Q the structure of a (trivial, principal) *fiber bundle*, with G the *fiber space* and R the *base space*.

lar media [Hat+13; Dai+; McI+16].

The meaning of the cost encoded by the metric $\mathcal{M}$ depends on the system physics, but at a high level it can be considered as the time it will take the system to execute the motion given a unit power budget. For systems moving in dry-friction environments, s can be specifically taken the energy dissipated while executing the motion [Dai+]; for the viscous friction model we consider in this paper, s is the time-integral of the square root of power dissipated [HDC17].

2.2.1 Extremal and optimal gaits

Locomoting systems typically move by repeatedly executing *gaits* — cyclic changes in shape that produce characteristic net displacements in position. Such cycles can be chained together to produce larger motions through the world.

Geometrically, a gait θ is a cyclic trajectory through the shape space with period T,

$$\theta : [0, T] \rightarrow R \tag{2.3}$$

$$\theta(0) = \theta(T), \tag{2.4}$$

and the system shape at any time t while executing the gait is $r = \theta(t)$.

Under the locomotion model in equation 2.1, the net displacement over one cycle of a gait is equal to the path integral of the local connection $\mathbf{A}$ over that trajectory. By an extension of Stokes' theorem, this displacement can be approximated[2] as the integral of the *curvature* of $\mathbf{A}$ over a surface θ_a bounded by the gait,

---

[2] The quality of this approximation depends on the choice of body frame for the system, which can be optimally selected once $\mathbf{A}$ is calculated in an arbitrary convenient frame. See [HC11b; HC13; HC15] for further discussion of this point.

$$g_\theta = \oint_\theta g\mathbf{A}(r)\,dr \approx \iint_{\theta_a} \overbrace{\mathbf{dA} + \sum\Big[\mathbf{A}_i, \mathbf{A}_{j>i}\Big]}^{\text{curvature } D\mathbf{A}}. \tag{2.5}$$

The curvature $D\mathbf{A}$ (formally, the total Lie bracket or covariant exterior derivative of $\mathbf{A}$ [HC15]) measures how much the coupling between shape and position motions changes across the cycle, and thus how much displacement the system can extract from a cyclic motion. Its components $\mathbf{dA}$ and $[\mathbf{A}_i, \mathbf{A}_j]$ are the exterior derivative (curl) and local Lie bracket of the system constraints, and respectively capture the net forward-minus-backward motion and parallel-parking motion available to the system. They are calculated as

$$\mathbf{dA} = \sum_{j>i} \left(\frac{\partial \mathbf{A}_j}{\partial r_i} - \frac{\partial \mathbf{A}_i}{\partial r_j} \right) dr_i \wedge dr_j \tag{2.6}$$

and

$$[\mathbf{A}_i, \mathbf{A}_j] = g^{-1} \left(\frac{\partial (g\mathbf{A}_j)}{\partial g} \mathbf{A}_i - \frac{\partial (g\mathbf{A}_i)}{\partial g} \mathbf{A}_j \right) dr_i \wedge dr_j \tag{2.7}$$

$$= \begin{bmatrix} \mathbf{A}_i^y \mathbf{A}_j^\theta - \mathbf{A}_j^y \mathbf{A}_i^\theta \\ \mathbf{A}_j^x \mathbf{A}_i^\theta - \mathbf{A}_i^x \mathbf{A}_j^\theta \\ 0 \end{bmatrix} dr_i \wedge dr_j, \tag{2.8}$$

where the wedge product $dr_i \wedge dr_j$ is the basis area spanned by the ith and jth basis vectors.

For systems with two shape variables, $\mathbf{dA}$ and $[\mathbf{A}_i, \mathbf{A}_j]$ have only a single component (on the $dr_1 \wedge dr_2$ plane), and equation 2.5 reduces to a simple area integral whose integrand is the magnitude of $D\mathbf{A}$. Extremal gaits for these systems (maximizing net displacement per cycle) lie along zero-contours of $D\mathbf{A}$, maximizing the sign-definite region they enclose.

As a general rule, these extremal gaits are more interesting mathematically than as a motion for a robot or animal to follow. With the exception of sports such as basketball that explicitly count steps, displacement-per-cycle is not a useful quantity to optimize, as it leads to wasting time or energy eking out all the available displacement in the cycle, instead of executing smaller but more productive cycles more times. When considering the optimality of a gait, it is thus typically more useful to measure its efficiency by dividing the displacement the gait induces over each cycle by the effort or time required to execute it.

In our model, we take the efficiency γ as the ratio between net displacement g_θ it induces and the path-length cost s calculated in equation 2.2, $\gamma := \frac{g_\theta}{s}$. The path length cost, s, which in the viscous case determines the energy dissipated over a cycle under optimal pacing, can be computed independently of the pacing of the gait. This enables γ to represent the proper notion of efficiency for such systems [HDC17]. Note that maximizing this efficiency is equivalent to maximizing speed at a given power (or minimizing power for a given desired speed), and so gaits with this property are always the most desirable for effective locomotion, even when the goal is "move fast" instead of "move efficiently."

As discussed in [HDC17], optimally-efficient gaits are contracted versions of extremal gaits: they give up low-yield regions of $D\mathbf{A}$ in exchange for a shorter path length, and thus a smaller expenditure of power or time. These gaits lie along curves where the gradient of efficiency with respect to changes in the gait parameters p,

$$\nabla_p \gamma = \nabla_p \frac{g_\theta}{s} = \frac{1}{s}\nabla_p g_\theta - \frac{g_\theta}{s^2}\nabla_p s \tag{2.9}$$

is equal to zero. As further discussed in [RH16; RH17], the gradient terms in equation 2.9 can be expanded in terms of $D\mathbf{A}$, $\mathcal{M}$ and $\nabla\mathcal{M}$ evaluated along the gait. Given these expansions, the dynamics of optimizing γ resemble those seen in a soap bubble, with an inflating pressure provided by $D\mathbf{A}$ balanced against a surface tension corresponding to s.

2.2.2 Empirical geometric models

The geometric approach described above was originally developed for systems whose equations of motion can be shown from first principles to have the form in equation 2.1. Building on these results, we demonstrated [Hat+13; Dai+] that the constraint curvature $D\mathbf{A}$ is also a useful tool for understanding the locomotion of systems whose dynamics are less "clean," and are only tractable through numerical modeling or empirical observation.

In these previous works, we first used nonlinear models [Hat+13] or experimental measurements [Dai+] to sample the relationship between $\mathring{g}$ and $\dot{r}$ across the tangent bundle TR. We fit a linear form to this relationship on a grid of tangent space base-points T_rR, giving $\mathbf{A}$ on a sampling of the shape space, from which we then calculated the components of $D\mathbf{A}$ as per equation 2.6 and equation 2.7. Plotting the curvature over the shape space then allowed us to directly identify effective gaits for translation and rotation for three-link and serpenoid system geometries, following the procedure illustrated in Figure 2.1.

2.3 Oscillators and Data-Driven Modeling

The second thread of prior work that this paper draws upon is a robust theory of gaits as oscillators, combined with a statistical approach to data-driven model construction. For an observed physical system, it is not always known a priori what this limit cycle is, what the dynamics of attraction to the limit cycle are, or even what the precise period of oscillation is. These properties of the gait can be extracted using techniques of data-driven Floquet analysis (DDFA)[Rev09; RK15a], the key elements of which we review below.

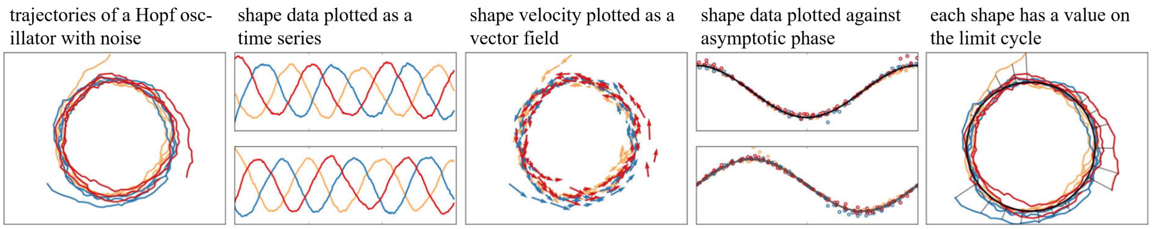

Figure 2.2: Data-driven Floquet analysis applied to a Hopf oscillator with system noise. Trajectories of the oscillator converge to a (noisy) cycle (left plots, three colors, one per trajectory). This cycle appears as a circle in state-space (extreme left) and as sinusoidal time series (second from left). By differentiation, we obtain vector field samples at the data points (middle). We estimate the limit cycle as a function of phase (second from right) computed using the phase estimator from [RG11] providing a canonical map from every trajectory point to a point with identical phase on the limit cycle (thin black lines, right plots). Such surfaces of constant phase — *isochrons* — form radial lines in the (particularly simple) case of the Hopf oscillator.

For simplicity exposition, we will assume all observations come in a single regularly sampled time-series consisting of (g_n, r_n) position and shape samples, which can be numerically differentiated (e.g., with a second-order Kalman smoother [RTS65; RG]) to augment the samples with velocities $\dot{g}_n$, $\dot{r}_n$, and $\overset{\circ}{g}_n = g_n^{-1}\dot{g}_n$. From oscillator theory [GH83; RK15a] we know that every exponentially stable oscillator (which we assume this to be) can be parameterized with a phase coordinate $\varphi : R \to [0, T) \subset \mathbb{R}$ based on the following rules:

1. Each point on the limit cycle has a unique phase value, spaced such that limit cycle trajectories advance in phase at rate $\dot{\varphi} = 1$.
2. Each point not on the limit cycle inherits its phase value from a corresponding point on the limit cycle, selected such that trajectories starting at the two points ultimately converge. The set of all points sharing a value of φ are called an *isochron* of the oscillator, and the trajectories of the oscillator advance across isochrons such that $\dot{\varphi} = 1$ everywhere.

Our modeling process was as follows: we assigned each sample n a phase φ_n via a *phase estimator* such as that presented in [RG08], which takes the multivariate time-series of measured data from an oscillator and gives a phase estimate for every data point. Figure 2.2 provides a visual example.

Once we grouped the samples by phase, we modeled the limit cycle (nominal-gait-as-executed) by computing a pair of Fourier series θ_0 and ω with respect to the phase: $\theta_0(\varphi_n) \approx r_n$ was fitted to the sampled shapes, and $\omega(\varphi_n) \approx \dot{r}_n$ was fitted to the shape velocities. Because each of θ_0 and ω is computed from its own noisy dataset, the condition $\dot{\theta}_0 = \omega$ need not be satisfied after this fitting procedure. We create a self-consistent model θ of the limit cycle by producing the analytical integral of ω, and using a matched filter to combine this integral with the θ_0 estimate to obtain a single self-consistent cyclic trajectory. Past experience [Rev09] has shown that this estimation procedure provides a better representation of the limit cycle than the directly fitted shape model θ_0.

2.4 Data-Driven Modeling of the Connection

The gait analysis methods described in §2.2 provide a powerful link between gaits' optimality and their geometry. Their utility, however, depends on having a model for how small shape changes induce body motion changes. For systems that experience complex interactions with their environments, such models are not readily available from first principles (even if their net effect can be modeled as the linear relationship in (2.1)), and exhaustive empirical evaluations [Dai+] become infeasible as we move to system with many shape variables and/or limited control affordances.

Conversely, the data-driven methods described in §2.3 are able to extract a meaningful model of a system from noisy measurements. This model, however, is limited to the specific gait being executed and does not provide context for comparing the gait against other motions the system could execute, or for optimizing the motion.

Our key innovation in this paper is based on the observation that the data-driven modeling approach can allow us to quickly build up a first-order model of the connection in a tube around a given gait cycle. In turn, the first-order model allows us to rapidly compute the influence of any gait change within the model's domain of validity. Such computations allow us to numerically approximate, at this given gait cycle, the gradient at of any goal function computed from a gait with respect to any parameterization of gaits — even when this parameterization is fairly high dimensional and requires a great many "simulations" of gaits.

In this innovation, we exploit two properties of the geometric model: [1] the variational optimizer/definition of optimality described in equation 2.9 only needs to know $D\mathbf{A}$ along the given gait to identify the direction in which that gait can be perturbed to best improve performance. [2] $D\mathbf{A}$, being a two-form and thus a linear map, can be reconstructed at every point along a gait cycle using regressions applied to the relationship between g and r collected from experiments.

2.4.1 Analytic approximation of the connection near a gait

In this section, we introduce an approximation of the mechanical connection and the cost metric, both centered about a nominal gait. We then construct a procedure to estimate the local model elements from data. As discussed in §2.3, a gait cycle $\theta(\cdot)$ can be extracted from shape data r via data-driven Floquet analysis. Perturbations from this phase-averaged behavior are written as $\delta(t) := r(t) - \theta(t)$. These terms can be used to construct a first-order approximation of $\mathbf{A}(\cdot)$ in a neighborhood of the point-set $\operatorname{Im}\theta$ using its Taylor series,

$$\mathbf{A}_i^k(r)\dot{r}^i = \mathbf{A}_i^k(\theta+\delta)\dot{r}^i \approx \left[\mathbf{A}_i^k(\theta) + \frac{\partial \mathbf{A}_i^k}{\partial r_j}(\theta)\delta^j\right]\dot{r}^i, \tag{2.10}$$

where, as per Einstein index notation, $\mathbf{A}_i^k$ corresponds to the element in the k-th row and i-th column of $\mathbf{A}$. Including the derivative of the connection across the shape space allows us to estimate the connection for behaviors that aren't on the current gait cycle $\theta(\cdot)$, and that thus provide velocity samples in nearby, but not identical, tangent spaces of the shape space.

It is important to stress that $\frac{\partial \mathbf{A}^k}{\partial r}$ is not simply the Hessian matrix of g^k with respect to r around points on the gait, i.e. it is not the gradient of a gradient. The Hessian could only be computed if g were a function of r, but it is not. In fact, locomotion via gaits would be impossible if it were such a function since a cyclical change in r could not induce a net change in g. In particular, Hessians are symmetric operators, and the difference term that appears when calculating $\mathbf{dA}^k$ in equation 2.6 directly measures the system's ability to locomote along the k-th direction in terms of the asymmetry of $\frac{\partial \mathbf{A}^k}{\partial r}$. Similarly, the $[\mathbf{A}_i, \mathbf{A}_j]$ term from equation 2.7 measures the the covariant asymmetry of $\frac{\partial g \mathbf{A}}{\partial g}$ when the connection is expanded from local to global coordinates.

2.4.2 Estimating $\mathbf{A}(\theta)$ and $\mathbf{DA}(\theta)$ from data

Our input data was time series of the system shape r_n, shape velocity $\dot{r}_n$, and observed body velocity $\overset{\circ}{g}_n$, at sufficiently many time points $n = 1 \ldots N$. We begin our system identification process by applying the gait extraction algorithm described in §2.3, producing Fourier series models of $\theta(\cdot)$ and of $\dot{\theta}(\cdot)$. We then select M evenly spaced values of phase, $\varphi_1 \ldots \varphi_M$, to obtain $\theta_m := \theta(\varphi_m)$ and $\dot{\theta}_m := \dot{\theta}(\varphi_m)$ — the shapes and shape velocities of a system that is following the gait cycle precisely. We use these as the points at which we estimate the connection and its derivative.

For each cycle point θ_m we collect all shapes r_n that are sufficiently close, i.e. n such that $\|r_n - \theta_m\| < \delta_{\max}$. For notational simplicity, when both index n and index m appear in an equation below, we take the values of n to be restricted to only those sufficiently close time series points. We now define the offset between the shape sample and its current on-gait reference point as $\delta_n := r_n - \theta_m$.

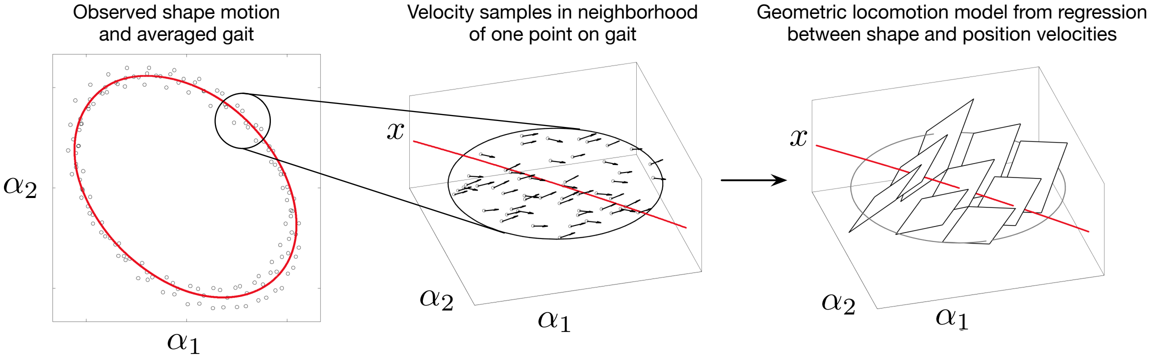

Figure 2.3: Illustration of the connection estimation process. We take the rhythmic data, group it by phase, and average using a Fourier series to obtain a periodic gait (left; red cycle). We collect shape velocity and body velocity data (middle; rooted arrows) within the neighborhood of point on the gait (black oval, left; zoomed in area, middle). Using these data we fit a first-order approximation of the connection model (black planes; right). We repeat this process for a collection of points on the gait cycle at fixed phase intervals, and fit the parameters of the estimated models with a Fourier Series to obtain a model of the connection that smoothly varies with phase. Further detail in §2.4.2.

Within each θ_m neighborhood, we now estimate the local connection and its derivatives by using a linear regression to find the slopes of the relationship between $\overset{\circ}{g}$, $\dot{r}$, and δ. Naively, this regression is the solution to the Generalized Linear Model formed by placing the Taylor-series expansion of $\mathbf{A}$ from equation 2.10 into the locomotion model from equation 2.1:

$$\overset{\circ}{g}{}^k_n \sim \left(\mathbf{A}^k_i\right) \dot{r}^i_n + \left(\frac{\partial \mathbf{A}^k_i}{\partial r_j}\right) \delta^j_n \dot{r}^i_n, \tag{2.11}$$

where $\left(\mathbf{A}^k_i\right)$ are the M separate estimates of $\mathbf{A}^k_i(\theta_m)$ and $\left(\frac{\partial \mathbf{A}^k_i}{\partial r_j}\right)$ are the M separate estimates of $\frac{\partial \mathbf{A}^k_i}{\partial r_j}(\theta_m)$.

When applied to samples generated from an oscillator as illustrated in Figure 2.2, this straightforward regression is biased by the shape velocity samples being centered around $\dot{r} = \dot{\theta}_m$ rather than $\dot{r} = 0$. We correct for this bias by re-centering the regression around $\mathbf{A}(\theta_m)\dot{\theta}_m$. We separate the perturbations of the shape velocity away from the gait cycle velocity from the influence of the gait cycle velocity itself by defining $\dot{\delta}_n := \dot{r}_n - \dot{\theta}_m$, and re-writing the GLM of equation 2.11 as (for velocity component k and each value of m):

$$\mathring{g}_n^k \sim \mathbf{C}^k + \mathbf{B}_j^k \delta_n^j + \left(\mathbf{A}_i^k\right) \dot{\delta}_n^i + \left(\frac{\partial \mathbf{A}_i^k}{\partial r_j}\right) \delta_n^j \dot{\delta}_n^i \tag{2.12}$$

where $\mathbf{C}^k := \mathbf{A}_i^k \dot{\theta}^i$ is the connection applied to the (unmodified) gait cycle shape velocity, and $\mathbf{B}_j^k := \frac{\partial \mathbf{A}_i^k}{\partial r_j} \dot{\theta}^i$ is the interaction effect of shape offset and shape velocity applied to the (unmodified) gait cycle shape velocity. Here $\mathbf{C}^k$ is a constant (with k, m fixed); and $\mathbf{B}^k$ is a ("co-")vector that acts on shape offsets from the gait, rather than on conventional tangent vectors. The $\left(\mathbf{A}_i^k\right)$ element is a true co-vector that acts on velocity offsets away from the typical gait velocity, and $\left(\frac{\partial \mathbf{A}_i^k}{\partial r_j}\right)$ is the interaction matrix of shape offsets and shape velocity offsets, offset being taken relative to the nominal gait cycle θ.

We compute the regression by writing it in matrix form and thereby posing the least-squares problem (for each k and m; indices k and m elided below for clarity):

$$\begin{bmatrix} \mathring{g}_1 \\ \vdots \\ \mathring{g}_N \end{bmatrix} = \begin{bmatrix} 1, & \delta_1, & \dot{\delta}_1, & \dot{\delta}_1 \otimes \delta_1 \\ \vdots & \vdots & \vdots & \vdots \\ 1, & \delta_N, & \dot{\delta}_N, & \dot{\delta}_N \otimes \delta_N \end{bmatrix} \cdot \begin{bmatrix} \widehat{\mathbf{C}}, & \widehat{\mathbf{B}_j}, & \widehat{\mathbf{A}_i}, & \widehat{\frac{\partial \mathbf{A}_i}{\partial r_j}} \end{bmatrix}^{\mathsf{T}} \tag{2.13}$$

where $\widehat{\ }$ indicates "estimated" and $\otimes$ is the outer product. For a d dimensional shape space, the row of unknowns on the right consists of $1 + d + d + d^2$ elements.

Once we have the model for every m, we construct a Fourier series model of each of the matrices of the GLM, allowing them to be smoothly interpolated at any phase value.

2.4.3 Estimating the metric

In the same manner as we estimatee **A**, we can estimate a Riemannian effort-metric $\mathcal{M}$ on the shape space by recording the differential cost of motion $\dot{s}$ along with the system kinematics, and then fitting these costs to a linearized expansion of equation 2.2 taken at the points θ_m using the matching n indices,

$$\dot{s}_n^2 \sim \dot{r}_n{}^{\intercal} \left[\mathcal{M} + \left(\frac{\partial \mathcal{M}}{\partial r_j} \right) \delta_n^j \right] \dot{r}_n. \tag{2.14}$$

This regression suffers from the afore-mentioned bias stemming from the $\dot{r}$ values being centered around $\dot{\theta}$ instead of 0;, so we recenter it in a similar manner as in equation 2.13. Additionally, because $\mathcal{M}$ is a symmetric tensor, only $\binom{d}{2}$ elements need to be estimated, reducing by about half the number of quantities to estimate. Details of this regression calculation are in the Appendix of [BHR18].

2.4.4 Comparison of estimates to previous work

This process is analogous to the processes we described for empirically estimating **A** and its derivatives in [Hat+13; Dai+], but offers some distinct advantages.

In our previous work, the shape velocity samples to identify **A** at a point all had to be in the tangent space of that point. Here, we have relaxed that requirement by fitting to a linearized expansion of equation 2.1 instead of equation 2.1 itself.

Furthermore, since our regression here uses intrinsic noise in the system, it provides an estimate of the average behavior under noise. The average behavior of a system when noise is added depends also on the variance of the noise. In the analysis here we account for the actual noise present in the system, rather than treating it as mere measurement error of a deterministic system.

The presence of system noise and the form of the linearized expansion allow for collection of data over a singular repeated gait cycle, rather than collection over the whole shape space (as was done for the prior model estimation methods).

2.4.5 Assumptions for the modeling estimation:

We make the following assumptions for modeling: (1) the deterministic part of the system's time evolution is governed by a connection; (2) the dynamics are subject to sufficient IID (independent and identically distributed) system noise to allow them to be identified; (3) noise is sufficiently small to allow a distinct rhythmic motion to be observed and modeled as a limit cycle oscillator representing a gait. For gait optimization, we further assume that the system is fully actuated and able to follow (on average) any trajectories we command.

2.5 Performance of the data-driven models

To benchmark the accuracy of our data-driven geometric modeling process, we compared its prediction of the body velocity for a test system against three system models that had various levels of knowledge about the "true" system dynamics used in the simulation. The test system had a geometric locomotion model of the form in equation 2.1, and its shape trajectories were generated via a noisy oscillator like that illustrated in Figure 2.2.

2.5.1 Reference models

As described in §2.4, we used a data-driven process to construct a phase varying first-order model of $\mathbf{A}$ at points θ_m along our observed gait cycle. Each r_n data point from the (noisy) trial was associated with a corresponding (phase-matched) point θ_n on the gait cycle,[3] which allowed us to compare several different models for the body velocity:

---

[3] These phase-matched θ_n points can be individually computed for each r_n, and so are not restricted to the previously-sampled θ_m values. Similarly, the estimates of $\mathbf{A}$ and its derivative from §2.4.1 are computed as Fourier series, and can thus be interpolated to any θ_n.

1. The *ground truth model*

$$\mathring{g}_{G,n} = \mathbf{A}(r_n)\dot{r}_n, \tag{2.15}$$

   in which each $(r_n, \dot{r}_n)$ pair is passed directly to the simulator dynamics, giving

2. The fully *data-driven model*, where the regression estimates of the Taylor expansion of $\mathbf{A}$ are used to approximate $\mathbf{A}$ at points off of the gait cycle, and $\mathring{g}_{D,n}$ is given by (2.12), used with the quantities estimated from (2.13).

3. An *analytic model*

$$\mathring{g}_{A,n} = \mathbf{A}(\theta_n)\dot{r}_n + \frac{\partial \mathbf{A}}{\partial r}(\theta_n)\delta_n\dot{r}_n \tag{2.16}$$

   that uses a Taylor-series expansion of the simulator model computed at the same point as the data-driven model, without using any regression or simulation data. This model tests the correctness of the regression in the data-driven model.

4. A *template projection model*

$$\mathring{g}_{T,n} = \mathbf{A}(\theta_n)\dot{\theta}_n. \tag{2.17}$$

   that projects each $(r_n, \dot{r}_n)$ data point onto its corresponding $(\theta_n, \dot{\theta}_n)$ values for the gait cycle that was used to derive the data-driven model. This approximation tests how much additional information is gained from the higher order term in the Taylor expansion.

Note that the template approximation in equation 2.17 can be considered as the leading term of the analytical approximation (after separating $\dot{r}_n$ into $\dot{\theta}_n$ and $\dot{\delta}_n$ components), and that the partial-derivative terms in equation 2.12 and equation 2.16 contain the information required to predict the effect of modifying the gait limit cycle.

2.5.2 Simulation setup: swimming with system noise

For our baseline system model, we used a three-link Purcell swimmer [Pur76] modeled as described in [HC13]. This system moves through a viscous fluid with linear drag, which we take as having a 2 : 1 lateral/longitudinal ratio. To demonstrate the ease with which we can extend our approach to systems with higher-dimensional shape spaces, we also considered a nine-link swimmer. Both are pictured in Figure 2.4 part A.

To simulate the effects of noise in the shape dynamics (e.g., weak or imprecise shape control), we generated the shape trajectories from sample paths of a (Stratonovich) stochastic differential equation, injected into the shape space:

$$\begin{aligned} d\varphi &= 1\, dt + \eta \circ dW_\theta \\ d\delta &= -(\alpha\, \delta)\, dt + \eta \circ dW_\delta, \\ r(t) &:= \theta_{\text{REF}}(\varphi(t)) + \delta(t). \end{aligned} \tag{2.18}$$

where $\theta_{\text{REF}}(\cdot)$ was a reference motion we specified as a Fourier series; α was the coefficient of attraction bringing the system back to the reference gait cycle; and η was a noise magnifier for the Weiner processes dW driving both phase noise and shape noise.

For all simulations in this paper $\alpha = 0.05$ and $\eta = 0.025$, chosen based on the superficial similarity the noisy trajectory ensembles have to experimental data we have worked with.

2.5.3 Model accuracy results

To illustrate the performance of our data-driven models, we examined the differences between motion predicted by the models in §2.5.1 when the reference gait was the extremal gait maximizing motion in the x direction, known from [TH07; HC13].

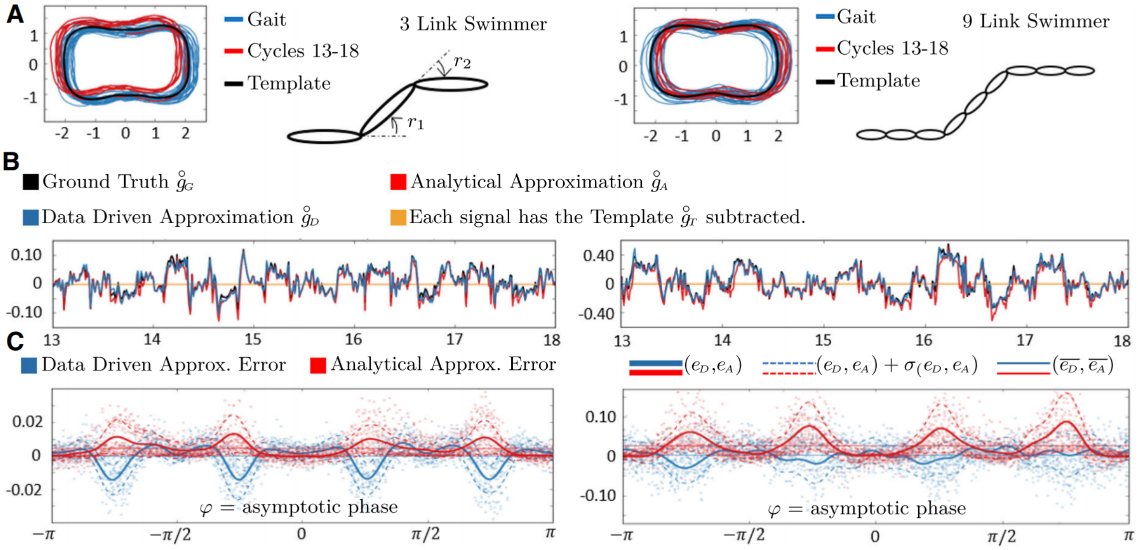

Figure 2.4: Comparison of model accuracy for 3 link and 9 link swimmers. [A] We drove each platform to follow the extremal gait for the three-link swimmer (black) generating 30 strokes (blue and red; plotted on first two principal components). Of these, we plotted Cycles 13-18 (red) in the time domain [B], showing the additional motion predicted beyond the template model by the ground truth model (black), the data-driven model (teal), and the analytic model (red). Because both analytic and data-driven models follow the ground truth closely, we also plotted a scatter plot of their errors as a function of phase [C], showing that the data-driven model (teal) has zero average error, unlike the analytic (red) model. As the number of DOF grows (right; 9 link plots) the mean (solid) and variance (dashed) of the data-driven model (teal) become smaller than those of the analytic model (red).

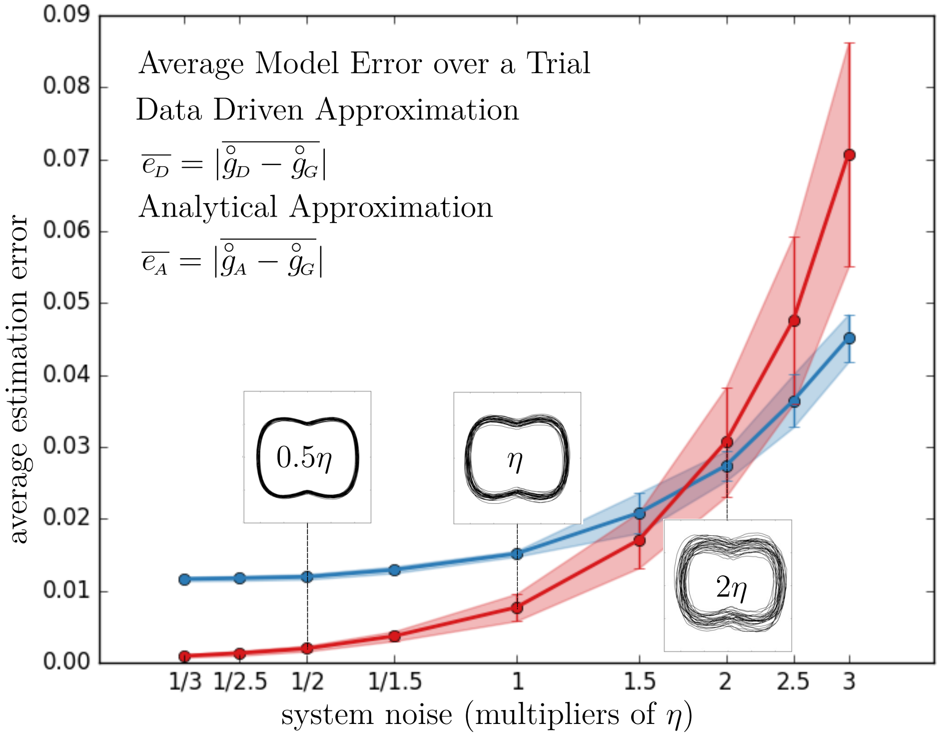

Figure 2.5: Comparing analytic and data-driven approximations. Given the same input gait and attraction laws of Figure 2.4, we plotted the accuracy of both models (data-driven in red; analytic in teal) over a range of system noise values (0.5η, η and 2η example trajectories in insets), and indicated the range of estimation error observed for x velocity over an ensemble of 20 trials at each noise level. System noise can be seen to strongly degrade the accuracy of the analytic model, whereas the data-driven model retains accuracy at high levels of noise, at the expense of accuracy at low noise levels.

We specifically chose an extremal gait as our example because non-extremal gaits should be even easier to model — perturbations around them have first-order effects. These results are shown in Figure 2.4.

In high noise regimes, the data-driven approach yields better models than the analytic Taylor expansion of the dynamics around the gait cycle. This effect is illustrated in Figure 2.5, which shows estimation error as a function of noise level for our example gait. The data-driven model also outperforms the analytic model when the system dynamics are very nonlinear, as for the nine-link swimmer at the right of Figure 2.4(C), and when the system noise is large.

These differences stem from the fact that the analytic model is a linearization of the system dynamics that extrapolates the system dynamics from their values on the gait cycle, whereas the data-driven approximation acts like a secant approximation to a curve, and averages the rate of change of the system dynamics across the neighborhood of the cycle. At the limit of large samples and small noise, the data-driven model approaches the analytic model. Thus, at the limit for many samples and finite noise, the data-driven model should always out-perform the analytic model — giving the best linearization for prediction over the available data, rather than the linearization locally at the gait cycle. However, with finite sample sizes the accuracy of the estimated linearization can suffer, allowing the analytic model to out-perform the data-driven one.

2.6 Data-Driven Geometric Gait Optimization

Given both the model of the connection (from §2.4.2), and the model of the cost metric (from §2.4.3) we can evaluate the efficiency of gait cycles in a Sobolev neighborhood of an initial gait cycle. In particular, this allows us to compute the gradient of efficiency (as in equation 2.9) and use a gradient ascent optimization scheme to optimize gait efficiency.

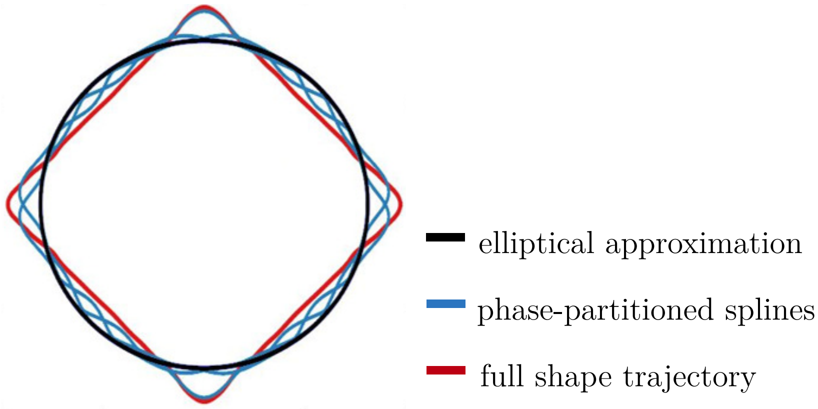

Figure 2.6: Illustration of gait parameterization as an ellipse with bump functions. In this parameterization each gait starts out as an ellipse – the image of a first-order Fourier series (black). To this we added a collection of overlapping, compactly supported, cosine window bumps. The number of bumps is the only order parameter for complexity of the model (here order 30). The sum of the circle and the plotted individual bumps (teal) combine together to give a diamond shaped gait cycle (red).

In implementing the gait optimization procedure two key choices to be made: (1) How is the space of gaits to be represented? (2) How big a step should the optimizer take along the gradient each time it is computed?

2.6.1 Gait parametrization

For our reference implementation of the gait optimization process, we constructed gaits in which the motion of each shape variable (here, joint angle) as the sum of a set of compactly supported bump functions added to first-order Fourier series. Each shape-space coordinate of the gait $\theta_i(t)$ is thus given, for an order N_o parameterization, as:

$$r_i(t) := c_i + a_i \sin(\Omega t + \phi_i) + \sum_{k=0}^{N_o} u_{i,k}\, w\left(t - k\frac{2\pi}{N_o}\right) \tag{2.19}$$

$$w(x) := \begin{cases} 1 + \cos(xN_o) & |xN_o| < \pi \\ 0 & |xN_o| \geq \pi \end{cases}, \tag{2.20}$$

with gait parameters

$$p_i = (c_i, a_i, \phi_i, \{u_{i,k}\}). \tag{2.21}$$

By construction, only 30 window functions from the sum in Equation (3.11) can be non-zero at any time in the gait — making the sum fast to compute, and restricting the influence of each $u_{i,k}$ to only $1/N_o$ of the gait cycle. The expressiveness of this representation in a two-dimensional shape space is illustrated in Figure 2.6.

2.6.2 Choosing a step size

Once we have identified an efficiency gradient vector $\hat{p} = \nabla_p \gamma$ on the parameter space, our optimizer must decide how large a step α to take along the gradient. This step size should be informed by the size of the neighborhood around the current gait which was sampled in the most recent trial, which is in turn determined by the level of system noise in the trial.

To compute the step size, we first measure the noise at each phase bin m as the covariance matrix of off-cycle displacements in the neighborhood of that bin, $\hat{C}_m := \mathbb{E}\left[\delta_n {\delta_n}^{\mathsf{T}}\right]$, which we can interpolate to any phase φ via a Fourier series. We then take:

- θ_0 as the current gait;
- p_0 as the parameters of this gait;
- $p_\alpha = p_0 + \alpha\hat{p}$ as the parameters reached by stepping along the efficiency gradient by α; and
- θ_α as the gait defined by these parameters,

and calculate the Mahalanobis distance with respect to the sampling noise [Mah36]

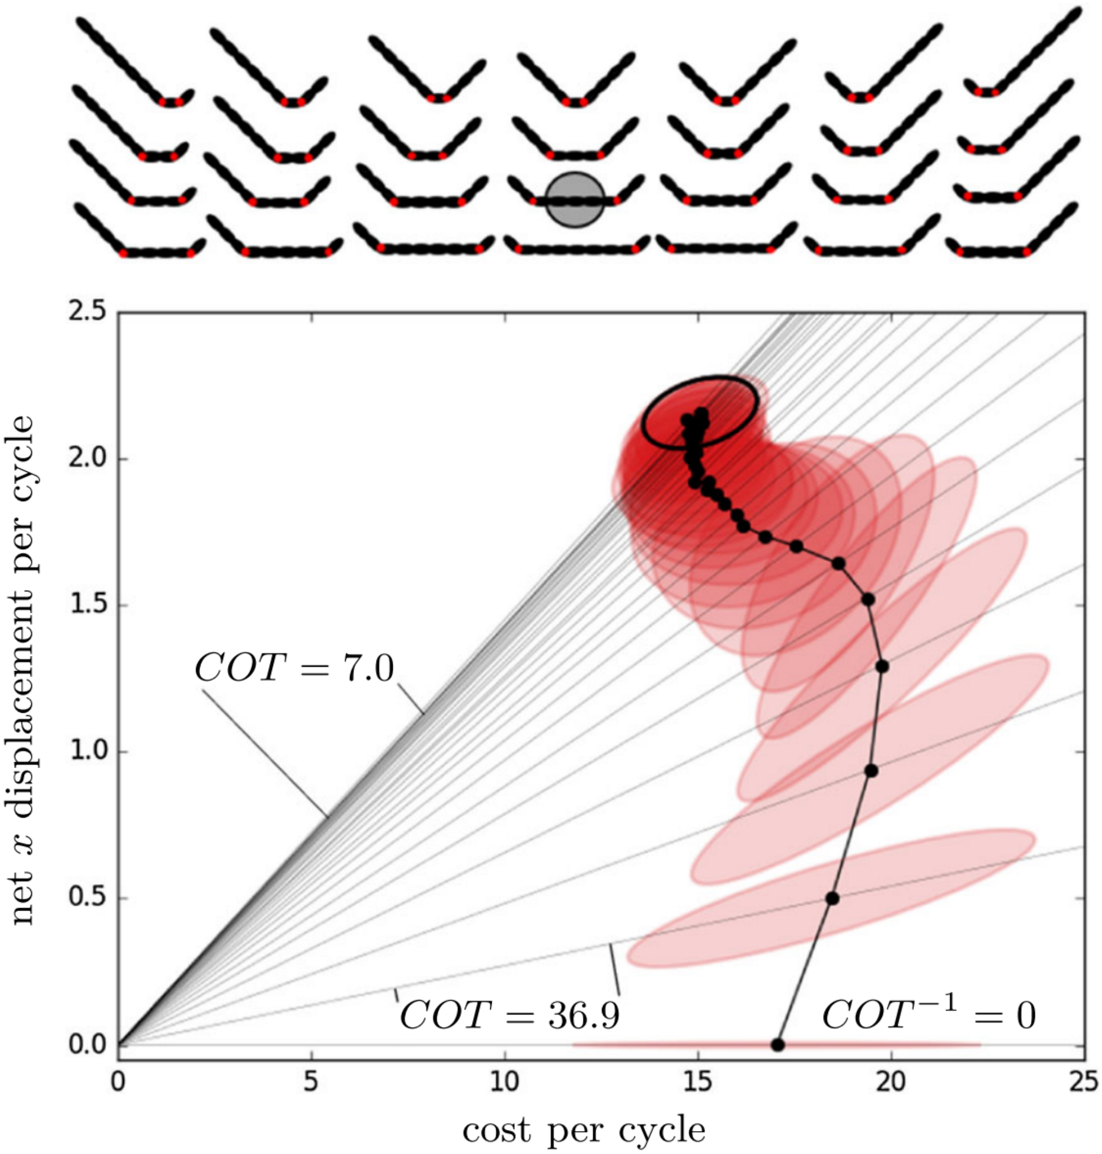

Figure 2.7: Optimization is insensitive to initial gait. We provided 28 different initial gaits (cartoons top) each with a different pair of joints (red dots in cartoon) following $r_i(\varphi) = \sin(\varphi)$, with all other joints set to constant angle 0. We optimized each initial gait 3 times, for a total of 84 optimization runs, and plotted the mean (black dots) and covariance ellipsoid (red) of the ensemble of gaits at every simulation iteration on axes of cost and displacement. In these axes, cost of transport (COT) corresponds to a slope. The initial gaits hardly move, giving a distribution along the horizontal axis, which improves to COT 36.9 after one iteration. As optimization progressed, all gaits moved toward the $COT = 7.0$ line, with the final (30th) iteration showing almost no progress and a fairly tight clustering of cost and displacement (black ellipse). Each optimization procedure converged to a serpenoidal motion, although these were not identical and retained some hint of the original choice of active joints. We used the initial gait highlighted (gray circle) for the noise regime testing in Figs. 2.8 and 2.9. By using reciprocal motions for initial gaits, we ensured (using the Scallop Theorem; see e.g. [Pur76; Lau11]) that all initial gaits have zero net displacement.

between the the updated and current gait cycles as

$$Z(\alpha) := \frac{1}{2\pi} \int_{\mathbb{S}^1} (\theta_\alpha - \theta_0)^\mathsf{T} C^{-1} (\theta_\alpha - \theta_0) \, d\varphi, \tag{2.22}$$

where θ_α,θ_0 and C are all functions of a phase parameter φ which we elide for clarity.

This distance measures the uncertainty of our data-driven modeling process and, generally speaking, grows with α. Presuming this growth to be monotonic, we can use a line search (e.g., a bisection search) to locate (to a user-selectable relative error tolerance; we used 5%) the α value at which $Z(\alpha)$ crosses some threshold value. For the experiments in the next section we used a threshold constant of 9.5.

The gait optimization framework can be summarized as a gradient ascent algorithm with careful considerations for the parametrization of the gait and step selection. Given an initial parametrization (detailed in §2.6.1), we collect experimental data (30 cycles in our results section) and compute the local motion-and-metric models. We extract a gradient on the efficiency of a motion with respect to the gait parameters by sampling many gaits in the neighborhood of the current policy, using the estimated local model to predict the performance of each sampled gait. We then determine the magnitude of the step size as described in §2.6.2. This allows the next gait parametrization to represent a behavior that is reliably informed by the data of the prior trial. Once the next gait is selected, we collect experimental data, repeating the above process. The termination criterion for the gradient ascent algorithm is a pre-specified number of iterations. A more advanced termination criterion will be explored in future work.

2.7 Swimming Gait Optimization Results

As a demonstration of our gait optimization framework, we applied our algorithm to a 9-link chain “swimmer”. All swimming behaviors shown were optimized with respect to the efficiency metric $\gamma = \frac{g_\theta}{s}$, which we report in units of body lengths per unit time at unit power. For any given power budget, this efficiency is inversely proportional to the mechanical cost of transport.

2.7.1 Optimization is robust to choice of initial condition

One important test of an optimization algorithm is its ability achieve good outcomes irrespective of initial conditions. To test this ability, we provided the system with gaits in which two selected joints follow identical sinusoidal inputs (no phase offset and amplitude of 1), the other joints attempt to hold at zero angle, and all joints are subjected to noise as discussed in equation 2.18. The power costs of these gaits depended on the lengths of the segments between the active joints, and, as illustrated in Figure 2.9, as reciprocal motions they produced no net displacement.

From each of these initial conditions, our optimizer consistently converged (within 30 trials at 30 cycles per trial) to gaits with a cost of transport of 7.0 ± 0.7. As illustrated in Figure 2.8, these resulting motions were very close to ellipses embedded in the eight-dimensional shape space, and produced serpenoid undulations traveling along the length of the swimmer. Qualitatively, the motions are in agreement with the conclusions about optimal swimming behavior in [SW89a], with the exception of maximizing the amplitude of the undulations at the mid-body of the swimmer. In this case, the amplitudes of the discovered gait are typically maximized near the joints that are excited in the initial gait. The reason for the discovery of this family of gaits and their relation to the global optimum in [SW89a] will be the subject of future work.

2.7.2 Robustness to noise level

A second test of optimizer performance, which is of particular importance to hardware-in-the-loop optimization, is its ability to tolerate a variety of noise levels

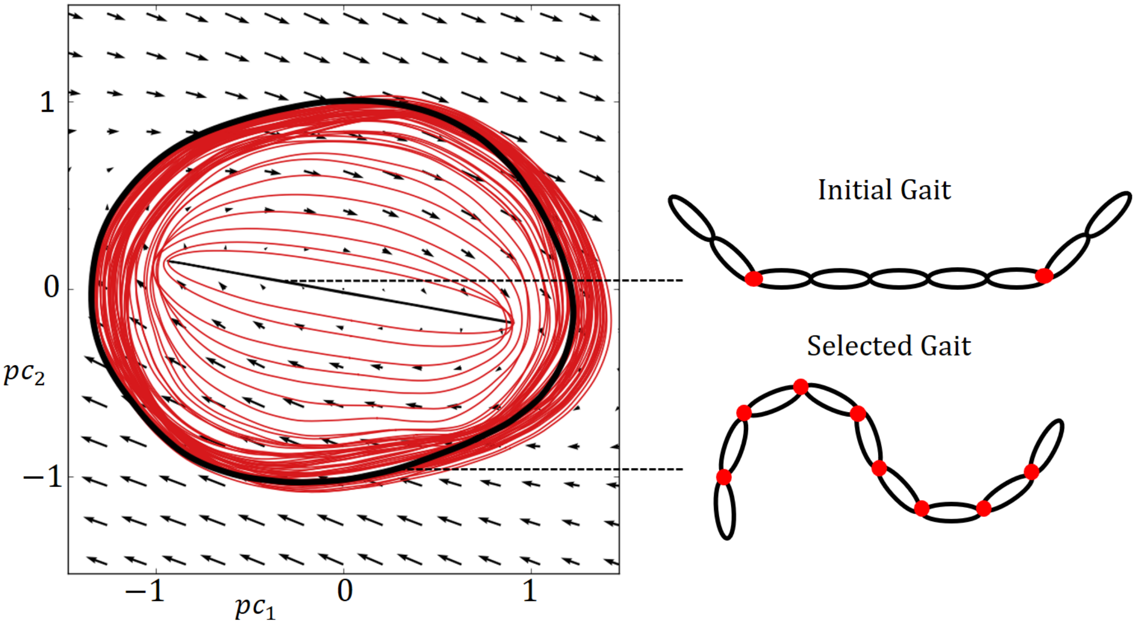

Figure 2.8: Visualization of gaits throughout an optimization. We projected all gaits onto the first two principal components of the final gait (viewed as embedded in $\mathbb{R}^8$) and plotted the projection of the x motion connection on that subspace (arrows). The initial gait (top cartoon), allowing only two joints to move (red dots in cartoon swimmer), is a line in the shape space coordinates (black line). The following iterations expand this contour as an ellipse and eventually embellish the ellipse with bumps (red closed ovals) leading to the final gait (black oval) and the serpenoid shape (bottom cartoon).

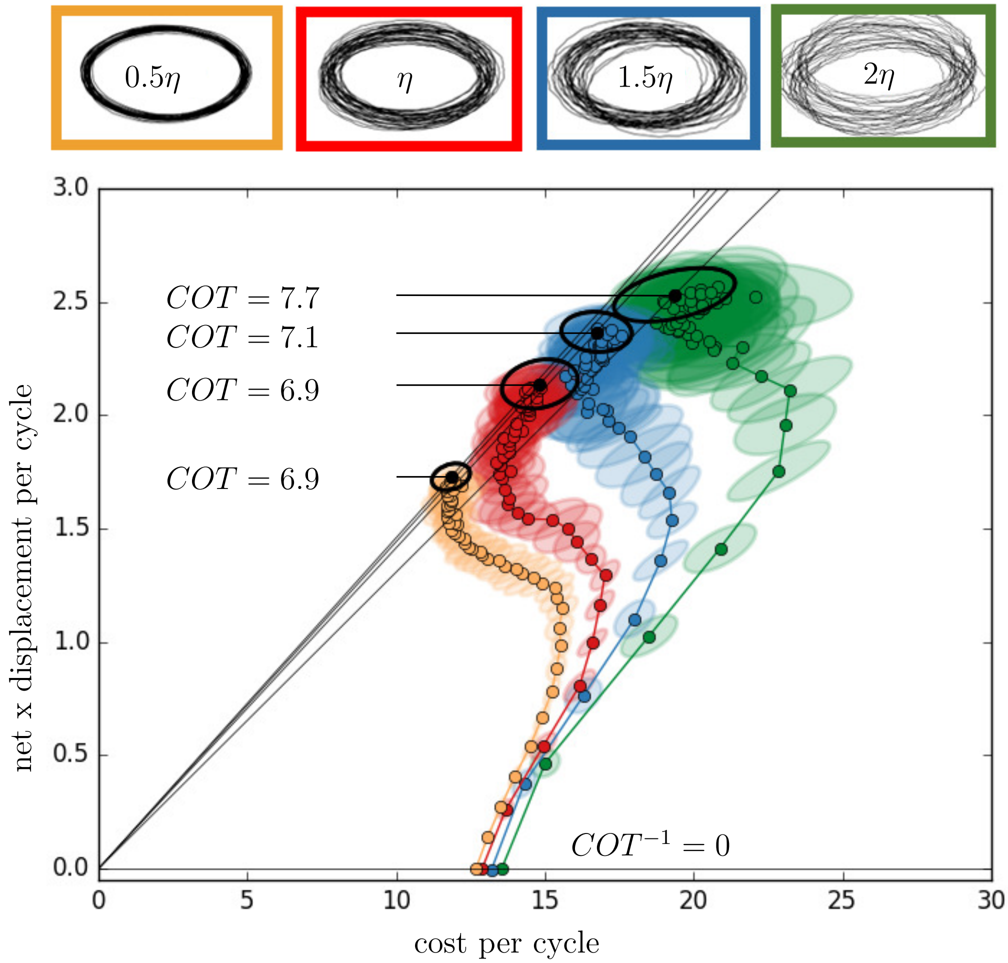

Figure 2.9: Course of optimization under different levels of noise. We started with the same initial gait (gray circle highlight in Figure 2.7 and top cartoon in Figure 2.8), but multiplied the noise level η of Eqn. (2.18) by $0.5, 1., 1.5, 2$ (colors yellow, red, teal, and green, respectively). For each noise level we plotted an example simulation to illustrate the noise level (ovals framed in color; top). We ran 48 optimizations at each noise level, allowing 60 iterations of 30 swimming cycles each, and plotted the mean (circle marker) and covariance (translucent ellipses) of these trials at every iteration of the algorithm, highlighting the final mean (black dot) and covariance (black ellipse). All gaits started unable to move, and reached COT 7.3 ± 0.4 with high-noise optimal gaits being slightly less efficient than low noise gaits (COT of mean 7.7 vs. 6.9). The two lower noise level achieved indistiguishable cost. It is notable that at higher noises, optimization moved away from the origin, producing larger motions with larger cost.

and produce comparably good results. To demonstrate this ability, we took a single starting gait (in which the active joints are each set two links in from the end, as illustrated in Figure 2.8) and optimized its motion under different levels of system noise.

For all four noise regimes tested, the system converged to serpenoidal motions with geometrically similar shapes (similar ratio of wavelength to amplitude), but with different numbers of waves along the body. As illustrated in Figure 2.8, the gaits found at different noise levels have similar costs of transport (with mean values ranging from 6.9 to 7.7), but the systems at higher noise levels tended towards gaits that were high-cost/high-displacement, at the expense of some efficiency.

Additionally, we note that at all noise levels, the systems initially modified their gait to increase their net displacement, then "pulled left" on the graph to reduce the cost of producing this displacement. The step sizes between trials are smaller on the low-noise systems, as they experience smaller perturbations during the trials, and thus have a lower bound on step size as discussed in§2.6.2.

2.8 Discussion and Conclusions

We have presented two main contributions: (1) a method for locally modeling a connection and a cost metric in the neighborhood of a gait cycle, based solely on the observation of noisy trajectories; (2) an algorithm for gait optimization that employs this method for gradient climbing.

Our modeling relied strongly on system noise to produce sufficient excitations to allow us to employ regression and identify the structure of the dynamics at every phase of the cycle. In this there is both a strength and a weakness. The strength comes from exploiting noise and being able to model systems with levels of noise comparable to those we have observed in animal and robot data. The weakness comes from relying on noise to be "system" noise – i.e. arising from true changes in the system state rather than from measurement errors. Measurement noise could mask some of the structure we expose by regression. It could also suggest to the optimization to move in a direction that is not achievable by the actual hardware. In experiments on hardware, constraints must be placed to avoid entering problematic configurations.

The great strength of our gait optimization algorithm is that it decouples the dimension of the gait parameter space from the dimension of the shape space and the number of trials needed. Once the model is identified for a gait, numerical evaluations of gait perturbations are very quick and allow the goal function to be differentiated with respect to hundreds of variables with little effort.

Some interesting features of the method emerged from the simulations run in [BHR18]. The nine-link Purcell swimmer was able to persistently converge to high efficiency swimming behaviors throughout various levels of process noise (injected perturbations into the shape space). This suggests that robustness to perturbations experienced may be a feature of the method. The stronger the experienced perturbations, the higher amplitude and lower efficiency the gait the swimmers converged to. It could be the case that the low amplitude, highest efficiency behaviors we observed were unavailable to the system under high process noise due to lower robustness qualities. This suggests that the process noise may dictate a robustness-efficiency tradeoff. Elaborating further on the relationship between noise level and which gaits are optimal may provide new insights into biological mechanisms of robust locomotion. Additionally, the swimmers converged to similar efficiency swimming behaviors from 28 separate initial gaits, all of which had an initial displacement of zero. This suggests that the modeling and optimization process was insensitive to the initial condition for the swimmers. An interesting research project might involve an investigation of the relationship between actuator redundancy and this property of insensitivity to the initial condition. In Chapter III, we explored the relationship between actuator redundancy and the ability to learn gait libraries for navigation under a variety of initial conditions as well as actuator malfunctions. In our investigation, we find that redundancy improves both the rate of learning gait libraries and the ability to recover from injury.

Some natural extensions of our work include expanding to a broader class of data-driven models outside those systems which admit connection-like models [Ost99; Blo+96; Baz+17]. One natural question which arises is that of systems that are "nearly" Stokesian – is there a useful and easy way to identify the notion of "nearly" Stokesian that translates to good predictive ability of the presented modeling tool? We investigate this for systems with momentum in Chapter IV §4.2. Remaining in the class of Stokesian systems, there is work to be done on how to extend this method to systems that are underactuated in the shape space. Some work has been done modeling and selecting gaits for systems that have Stokesian mechanics and elastic joints [Dea+20; RH20]. Of interest to us is the ability to model a general class of shape-underactuated, Stokesian systems from data. This interests us due to the ability to extend system identification and optimization techniques to the space of soft and compliant robots.

2.9 Optimization on Hardware

We tested the policy gradient optimizer on two experimental platforms. Like the swimmers, these robots consisted of a linked set of joints. To test the viability of the method, we ran the optimization process on a robot with four joints and the same robot with eight joints.

2.9.1 Methods for five-link wheeled snake

First, we implemented the optimization on a physical five-link wheeled snake robot. Under each link, a pair of wheels were co-aligned to decrease friction in the

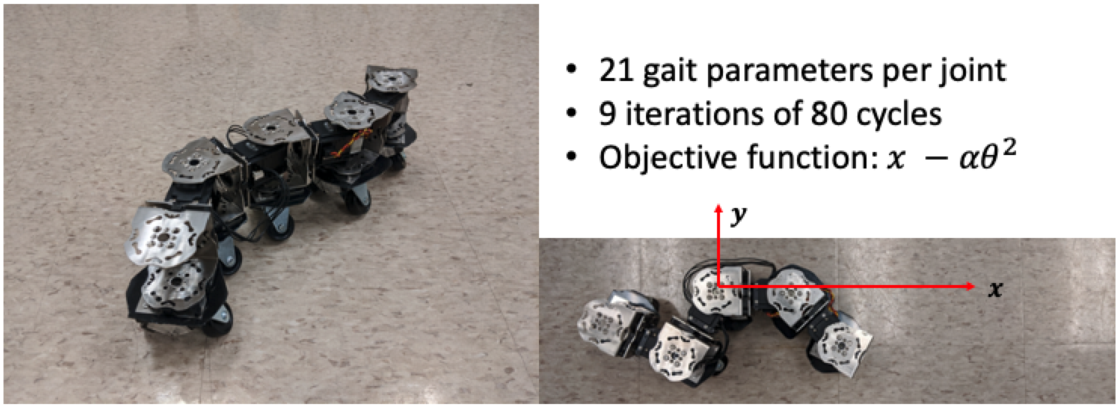

Figure 2.10: A pair of wheels under each module decreased friction along the x-axis of a given module, as defined by the coordinate axis on the middle module. The coordinate axis of the middle module defines the body x-axis for the entire system, with the y-axis as shown making up the lateral axis.

direction along the link. Figure 2.10 provides a picture of the system along with labeling of the body coordinate system. We captured the robot's position and orientation with three markers tracked with Qualisys Oqus motion capture cameras. We were able to pass joint trajectories to the Dynamixel RX-64 modules from a lab CPU (Intel Xeon CPU E3-1246 v3 running at 3.50GHz) over a CAT5 cable, which also supplied power to the system.

The parameterization for each of the four joints was an ellipse with 18 bump functions, using the representation motivated in §2.6.1. The corresponding optimization used 9 iterations of 80 cycles, with results reported in Figures 2.11 and 2.12. The objective function was to optimize displacement per cycle in the body frame with a quadratic penalty (scaled by α) on radial displacement per cycle (θ). To be specific, $\alpha = \frac{\pi}{L}$, where L was the body length of the robot (approximately 45cm).

2.9.2 Results for five-link wheeled snake

We seeded the geometric gait optimizer with a gait that oscillates a single joint and achieves no displacement per cycle. After 9 iterations of 80 cycles, we produced a

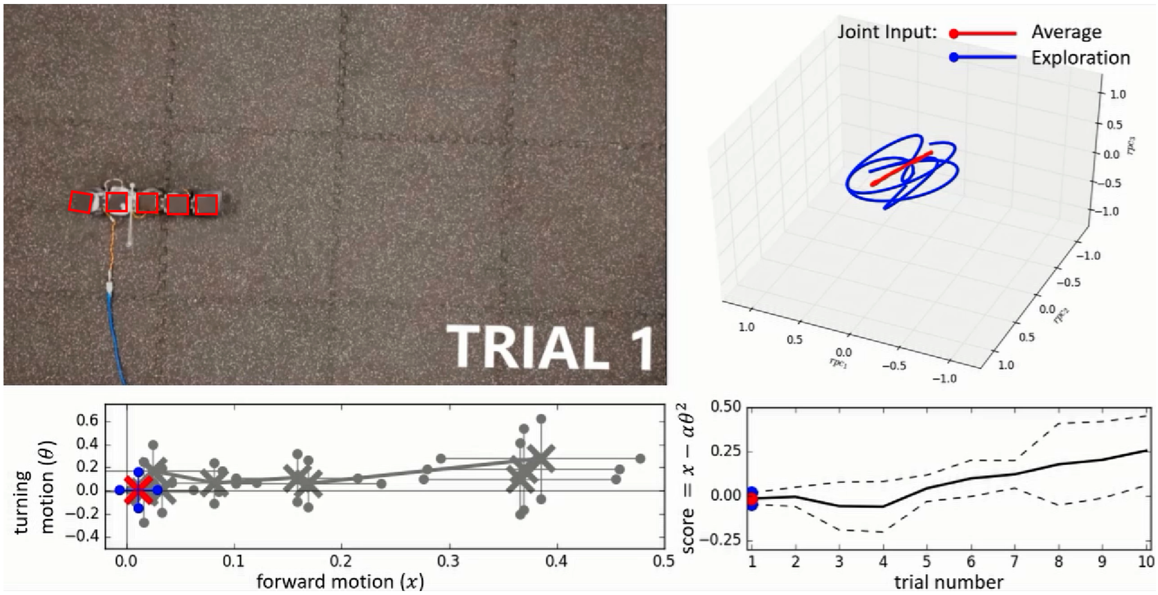

Figure 2.11: This figure documents the first trial on the five-link snake hardware in the loop optimization. The first gait provided to the optimizer was to oscillate the tail joint with an amplitude of 1 radian, driven at $\frac{1}{2}Hz$. The robot barely moved (top left, initial and final position after 5 cycles image blended with alpha=50; initial module locations also as red squares). We plotted the mean of the forward displacement per cycle and orientation displacement per cycle (bottom left, with means shown as red 'x's and standard deviations as blue dots connected by a black straight line). We plotted the score of the objective function across the 80 cycles of each iteration (bottom right, with means shown as a black line and standard deviations as black dashes). We also plotted the history of the trajectories sampled on the system along the primary 3 components of the sampling space (top right, average shape trajectory shown in red and added perturbations in blue). We computed the principal components via singular value decomposition on the entire sampling set after the 9 iterations of 80 cycles across 4 joints.

motion that achieved 45% body length per cycle translation motion while optimizing over 84 parameters. Running the system at 0.5 Hz, the optimization took 24 minutes. Results can be viewed in Figures 2.11 and 2.12.

2.9.3 Methods for nine-link wheeled snake

We added four more links to the robot and removed the high friction mat. Other than that, the methods are the same as for §2.9.1. We made this change because the nine-link wheeled snake was not strong enough to control shape on the high friction

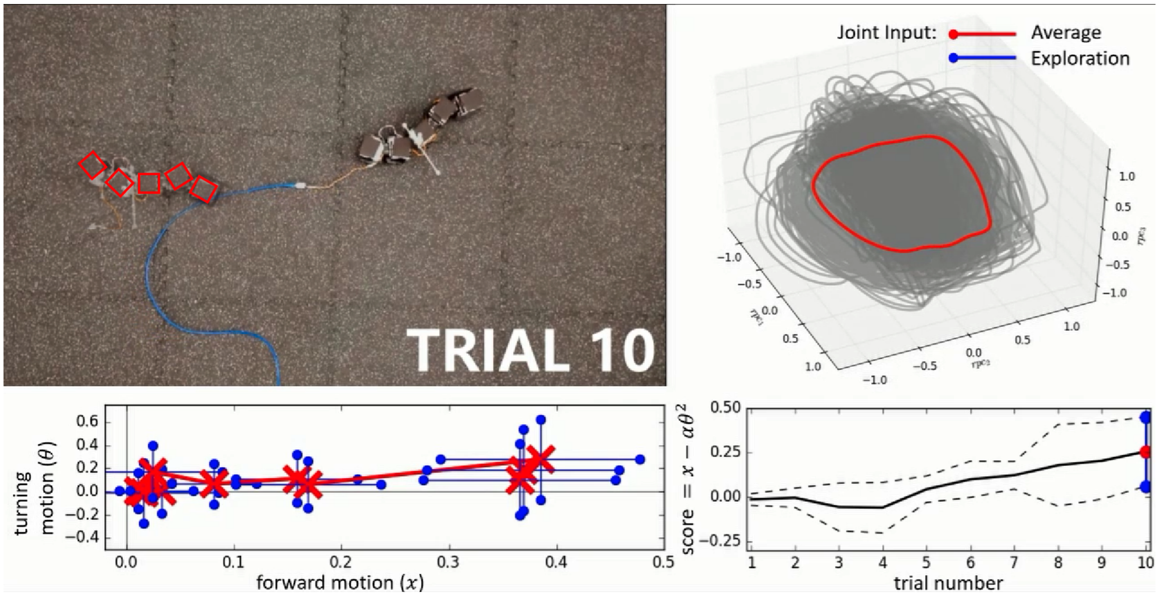

Figure 2.12: This figure documents the final trial on the five-link snake hardware in the loop optimization. We saw that after 9 iterations and 27 minutes of experimental data, the robot had found an effective strategy for locomotion on the high friction rubber mat. The meanings of the subplots are identical to those highlighted in Figure 2.11. While the optimization yielded a useful behavior, the cycles required to build the model took longer than those in simulation.

mat. Another challenge we observed for this system was a tendency to accumulate substantial lateral displacement over the course of a few cycles, so we placed a quadratic penalty on it in the subsequent optimization.

2.9.4 Results for nine-link wheeled snake

We found that our robot could build predictive models with just 30 cycles. Results can be viewed in Figures 2.13 and 2.14. After 360 cycles (12 trials of 30 cycles), the robot found an effective motion for translation. These results confirmed that the quality of results we saw on the Purcell swimmer could extend to relatively affordable and noisy hardware.

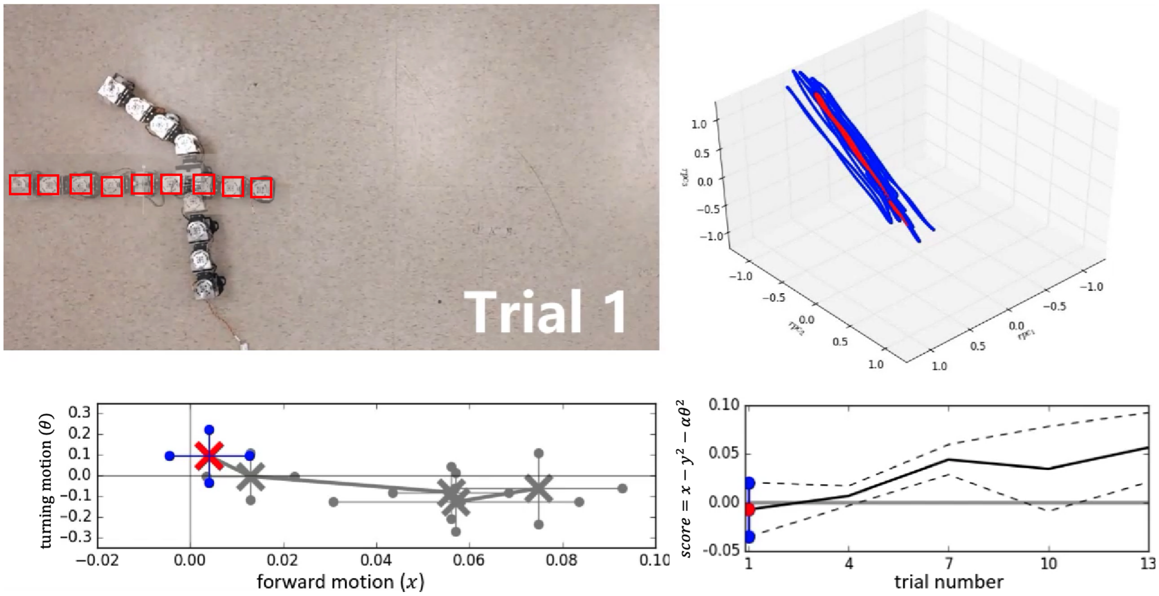

Figure 2.13: This figure documents the first trial on the nine-link snake hardware in the loop optimization. The first gait provided to the optimizer was to oscillate the third joint from the left with an amplitude of 1 radian, driven at $\frac{1}{2}Hz$. The robot barely moved (top left, initial and final position after 10 cycles image blended with alpha=50; initial module locations also as red squares). We plotted the mean of the forward displacement per cycle and orientation displacement per cycle (bottom left, with means shown as red 'x's and standard deviations as blue dots connected by a black straight line). We plotted the score of the objective function across the 30 cycles of each iteration (bottom right, with means shown as a black line and standard deviations as black dashes). We also plotted the history of the trajectories sampled on the system along the primary 3 components of the sampling space (top right, average shape trajectory shown in red and added perturbations in blue). We computed the principal components via singular value decomposition on the entire sampling set after the 12 iterations of 30 cycles across 8 joints.

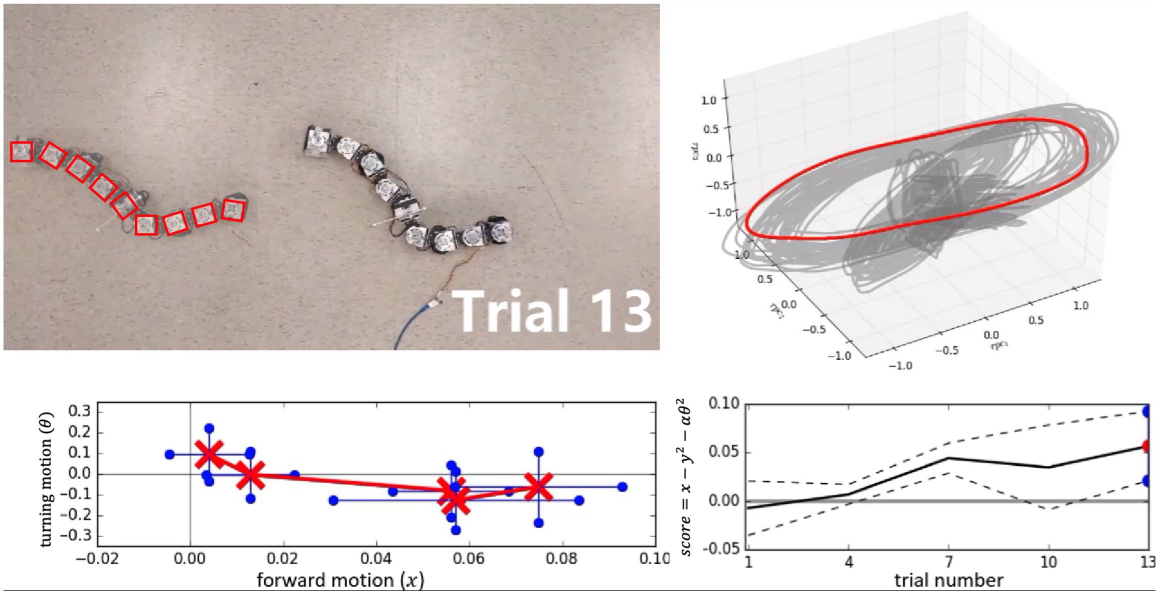

Figure 2.14: This figure documents the final trial on the nine-link snake hardware in the loop optimization. We saw that after 12 iterations and only 12 minutes of experimental data, the robot had found an effective strategy for locomotion on the laboratory floor. The meanings of the subplots are identical to those highlighted in Figure 2.13. The ability of a robot to optimize its behavior across 8 joints over 12 iterations of 30 cycles per iteration is comparable to the results we found in simulation.

2.9.5 Discussion on Hardware Results

The eight joint system built predictive models with less data than the four joint system. We think this was likely due to the smoother interaction between the wheels and the contact surface. For the five-link wheeled snake, the mat was creating stiction between the wheels and rubber, introducing nonsmooth estimates of the velocity of the robot by the motion capture system. Noisy velocity estimates may have been the phenomena that required us to sample more cycles to build models on the five-link wheeled snake.

CHAPTER III

Data-Driven Planning for Stokesian Systems

3.1 Motivation

One of the most common sub-problems in modern robotics is path-planning, and the choice of path is usually framed as a precise or approximate optimal control problem. When restricted to mobile robots moving through many practical environments, the path planning problem enjoys an additional important symmetry. Given the configuration of the robot body, the short-horizon movements it can execute are the same at nearly every point in space. This allows short time horizon *primitives* to be optimized offline and pre-cached, later to be composed sequentially to produce solutions to the full path planning problem. For example, a humanoid robot such as ATLAS can execute the same walking steps at any point on flat, unobstructed ground. To plan the motions of the robot walking through a building, one can sequence primitives for generating a collection of steps in the correct order instead of solving the full high-dimensional planning problem.

Unfortunately, the primitives often seen in such library-based plans are usually selected and created by hand. Primitives are often generated with constraints that help reduce the complexity of an individual planning problem. For example, a common choice for 2D motion, dating back to the turtle robots of the 1950s [Wal54], is to have linear translation and turning in place as primitives. However, this particular choice for generating movements is entirely arbitrary. A given robot may be far more efficient moving diagonally or turning while moving on an arc. The ability to optimize for a library of useful primitives can come to have critical importance when a robot is damaged, and the choice of available primitives might no longer correspond to any motion obvious to a human operator.

Here we present a method to optimize an entire primitive library concurrently so as to achieve the ability to efficiently plan over the space of body motions with that library. By optimizing for the *coverage* goal function we define, the library selected will be able to express desired short-horizon plans through composition of primitives from the library.

One approach for approximate optimal planning is to construct a state lattice [PKK09; KM09] – a discrete collection of states that can be generated by a library of primitives. Planning consists of sequencing primitives to travel along the lattice to approximate the total desired motion. Such previous work on state lattices suggests that a good collection of primitives are:

- *complete* – the space of desirable motions is densely populated
- *fast-to-compute* – the robot is able to select and generate primitives in real time
- *path optimal* – each individual primitive should be similar to a globally optimal path available between its start and end states.

When generating primitives, one has a variety of options to chose from [FMJ02; FDF05; Sch+05; Hau+08; PK11]. Strategies can include learning from demonstration as well as prioritizing spatial properties of the output trajectories of the system. Large primitive libraries are often winnowed down to save run time or increase the planning update rate. Our work can be viewed in part as a means for generating very small, very expressive libraries of primitives.

Our work can also be seen as a way to relax the standard assumption used in optimizing gaits, namely pre-specifying the direction [HC10; GGC] or turn rate of motion [Da+16; HDG17] over a single cycle. We observe that most of the value a primitive has is not intrinsic, but rather in its contribution to support other compositions of primitives available to a planner and the overall needs of the planning task. We thus provide a way to evaluate libraries of primitives rather than their individual characteristics. Primitives that have negligible exploration value in isolation may be critical to more densely maneuvering through space. We demonstrate how our coverage measure values such primitives rather than discards them.

Using our approach is nearly paradigmatically opposite to traditional behavior learning in robotics. We allow for the optimizer to "ask" the robot what ways are convenient to move, rather than dictating how the robot should move apriori. A subsequent advantage is that mechanical designers can rethink common design criteria for locomotors. Typically robots acting on a planar workspace are designed to have at least one mode by which gaits translate the system without rotating it. This preference may simply be the result of anthropocentric bias. It is how humans move to avoid disorientation and dizziness, but it is not a universal requirement for effective locomotion. The coverage measure, being devoid of such biases, allows a broader range of robot mechanisms to score highly. Crucially, it can also potentially allow broken robots to recover their ability to plan motions by rapidly regenerating a primitive library while damaged.

3.1.1 Overview of the sequel

Below we briefly review Lie groups in §3.2. This representation lent itself to the coordination of primitive libraries as a sequence of group actions acting on a Lie group of body locations. Using this representation, we defined coverage in §3.3 and provided examples of how it can be computed on the rigid body groups $\mathsf{SE}(2)$

and $\mathsf{SE}(3)$. In §3.4 we used this coverage to discuss a collection of toy systems whose locomotion ability becomes easy to appreciate through our approach. Following this, we translated this framework of primitive optimization to the world of gait driven systems in §3.5. There we paid special attention to highly damped systems, where the task of chaining primitives can be highly simplified. We presented coverage optimization of gait libraries for some Purcell swimmer models in §3.7. After this, we used coverage as a tool to investigate the ability of the Purcell swimmer to recover from joint locking in §3.8. Finally, we emphasized the ability of the optimization to work on unintuitive robots, even when we do not specify the robot kinematics, mass distribution, or material properties. We demonstrated this by the ability of a robot made of tree branches to gain the ability to navigate with less than eleven minutes of experimental data by optimizing for coverage.

3.2 Expressing Motion Through the Space of Discrete Actions

To represent motion, we assumed that the configuration space $\mathcal{Q}$ of our moving robot could be factored as a product of a shape R and a (generalized) position G. This generalized position is a Lie group, typically a sub-group of the rigid body motions $\mathsf{SE}(3)$. In this work, we restricted our attention to motion in the ground plane, $\mathsf{SE}(2)$. We produced motion using *gaits*, which we take to mean periodic changes in shape that produce a predictable body motion, i.e. a gait b is a function $\gamma_b : S^1 \to R$ that produces a body motion $M_b \in G$. For example, a stride consisting of a left step followed by a right step is a gait cycle of human walking. Given a finite selection of gaits and a means for switching between them, a planner can produce any motion that corresponds to a word in the group elements (letters) those gaits generate. For example, with gaits γ_a and γ_b, and provided any sequence is allowed,

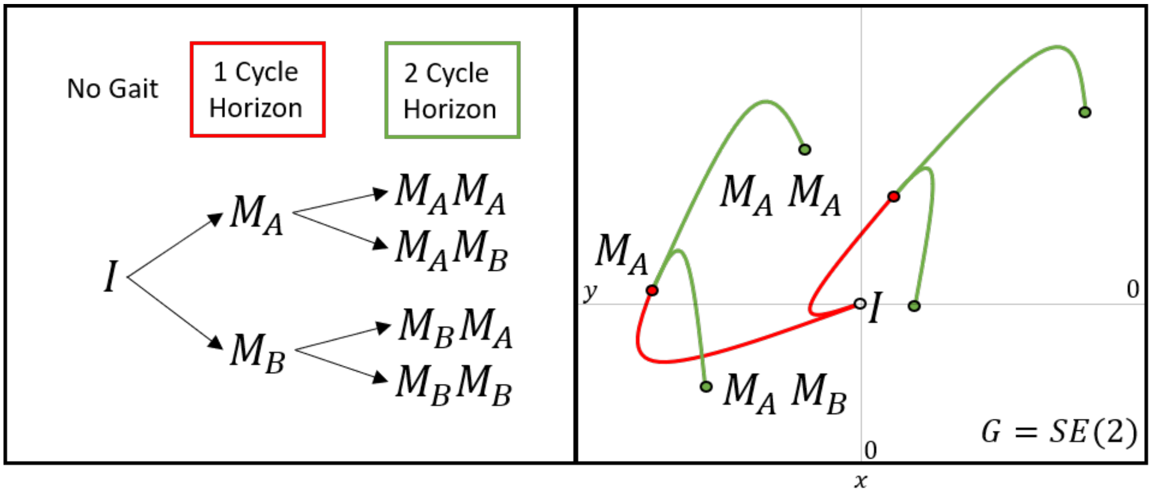

Figure 3.1: Illustration of composing gait cycles. Here, the two group actions (M_A, M_B) are applied in various orders and combinations. An n-step finite horizon planner considers words, a concatenation of group action letters, of length n. For a two letter action library, n step planners consider 2^n paths (trees in the left panel). We illustrated a possible case of such motions. By assuming that the robot is oriented tangent to the direction of motion, the resulting motions can be represented by their projection on the translational plane (right panel).

one can produce the motions $(I, M_a, M_b, M_a^2, M_a M_b, M_b^2, M_b M_a, M_a^3, \ldots)$. Figure 3.1 provides visualization of what this representation looks like for motion planning in a planar workspace.

In this paper, we restricted our discussion to primitive libraries consisting of single cycles of different gaits as the primitives. Here, we assumed the gaits are connected in internal state at their start and end configuration. This is not generally the case and requires careful treatment in §3.5.

3.3 Specifying the Loss Function

As conventionally practiced, a motion planner is given some parameters x, a means to generate motions $M(x) \in G$, and some loss function which it will minimize. Because we focused on reachability, we took the loss function

$$\tilde{\eta} : G \to \mathbb{R}^+ \tag{3.1}$$

to be purely a function of endpoint. Including additional factors in the loss functions for individual primitives is only a matter of book-keeping, provided the loss function of the overall path is additive in those of its constituent primitives. We assumed that the loss function is written relative to some desired goal position G of the motion, and defined a relative (local) loss function using the Lie algebra $\eta(\xi) := \tilde{\eta}(\exp(\xi)G)$. We then optimized for the parameters x of the primitive with respect to the loss function $x \mapsto \eta \circ \log(M(x)G^{-1})$. Any left invariant distance metric for $\mathsf{SE}(2)$ or $\mathsf{SE}(3)$ provides practical implementation of $\tilde{\eta}$. Picking such a metric boils down to a choice of a constant that relates the loss of translation errors to the loss of orientation errors, and thus this choice is application specific. In this work, we chose this parameter to make a half rotation on any axis equal to a body length displacement.

We set up our optimization as follows: let $\mathbf{G} := \{G_i\}_{i=1}^n$ be a set of goal motions and $\mathbf{W} := \{w_i\}_{i=1}^n \subset \mathbb{R}^+$ be a corresponding set of weights. Let $\mathbf{M} := \{M_j\}_{j=1}^m$ be a set of achievable motions.

We defined the coverage cost

$$h(\mathbf{M}) := \sum_i w_i \min_j \eta \circ \log(M_j G_i^{-1}). \tag{3.2}$$

We further defined $h_k(\mathbf{M})$ as the cost of the set of words of k or fewer elements of $\mathbf{M}$. The coverage cost is the sum of the costs of the best approximations available for G_i, given the achievable M_j and weighted by the weights w_i for each G_i.

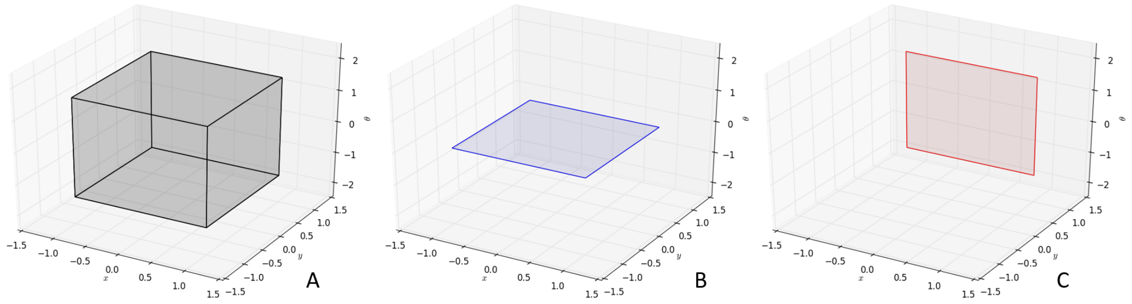

Figure 3.2: Expressive power of the coverage cost. One has a variety of choices for placement and weighting of coverage points. We provided some suggestions for various design goals on the space of planar rigid body motions. A user can prioritize versatility (panel A), zero-rotation translation (panel B), or right lateral movement (panel C). Volumes and planes are suggested regions for the user to evenly distribute uniformly weighted coverage points G_i.

3.3.1 Higher order maneuvers

One of the surprising insights of nonlinear control is that the non-commutativity of control actions can make reachable the iterated Lie brackets of a control distribution [SM92; Sas13]. The discrete primitive library equivalent of this insight is the observation that the *commutator* word $M_a M_b M_a^{-1} M_b^{-1}$ can at times reach directions that no word of the form $M_a^n M_b^m$ could reach. Thus, designing $h_k(\cdot)$ such that $k \geq 4$ allows these higher-order maneuvers to be included. It is, however, important to note that the coverage computation time scales exponentially with k. For this reason, we used $k = 4$ in our implementations here.

3.3.2 Design choices for coverage points

The coverage cost presented offers a user the ability to specify both the placement and weighting of coverage points. The selection of the points and weights can radically change the priorities of the optimizer. A user prioritizing versatility may want the robot to be able to reach all parts of its local position space. They can place a uniformly weighted set of points distributed evenly within some volume around the identity motion (Figure 3.2 A). Another user may wish to find a combination of gaits that translate while preserving orientation. That might correspond to a coverage point distribution in a thin wedge near the 2D slice of $\mathsf{SE}(3)$ corresponding to no rotation (Figure 3.2 B). Such maneuvers might be useful for an inspection robot that needs to maintain a visual field of view while moving. If one had a more specific navigational goal, *e.g.* finding a way to translate laterally to the right, such a goal can also be captured (Figure 3.2 C).

The w_i weighted collection of coverage goal points G_i can be seen as a discrete approximation to a measure on the group. Increasing the number of goal points in a region while keeping the total weight constant implies a preference for higher resolution in that region. Changing the weight while keeping the goal points unchanged implies an increase or decrease in the importance of approximating those goal motions with the primitive library.

3.4 Coverage Invites Non-traditional Mechanical Designs

Here we intentionally designed two robots that move in unconventional ways. The first cannot translate without rotating. The second has a trilateral symmetry. Often, roboticists do not consider such systems because their mobility is non-intuitive. Yet we have shown below that both systems can move quite effectively in $\mathsf{SE}(2)$.

3.4.1 Introducing two new mechanisms

Both of the mechanical system models we present are swimmers that operate at the limit of low Reynolds number fluid dynamics [Pur76], where friction dominates inertia. The motion of these systems is fully dominated by the drag forces induced by the internal velocities of the robots shape variables $r \in R$ and body velocities $\mathring{g} \in \mathsf{TSE}(2)$. The motions of these systems can be usefully inspected using the tools

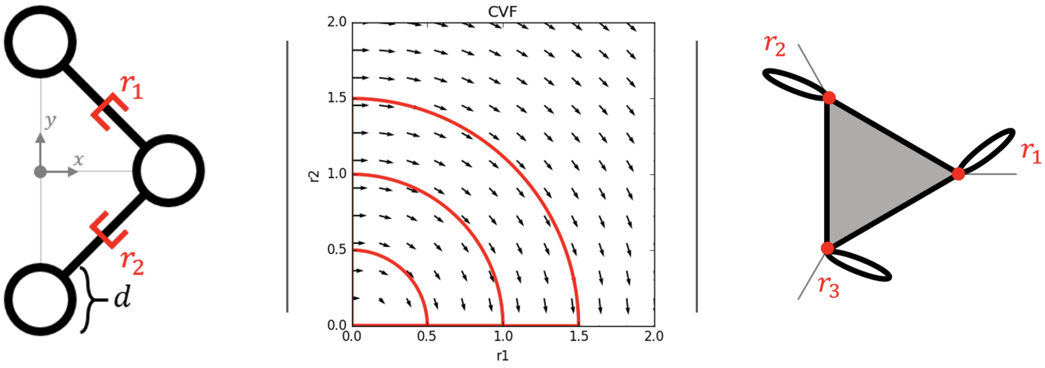

Figure 3.3: Here we describe two mechanical systems that may appear as unconventional travelers. The two-slider swimmer (left) can move spheres along prismatic joints. The motion simultaneously induces a thrust on the system while changing the geometry of drag forces acting on the system. We plotted the gaits selected for the two-slider swimmer on the rotational connection vector field [HC10] of the two-slider swimmer (middle). This provided insight into how shape change can influence body velocity. We can see that paths (shown in red) that start in the corner at the origin, travel along a shape axis, sweep at a constant radius to another axis, then return to the origin. The connection vector field aided gait selection of the two-slider swimmer, which is discussed in §3.4.2.1. The three-branch swimmer (right) has three links that can rotate, fixed to the end of a triangle. Since the shape space of the three-branch swimmer is not restricted to planar representations, we selected gaits in a different way.

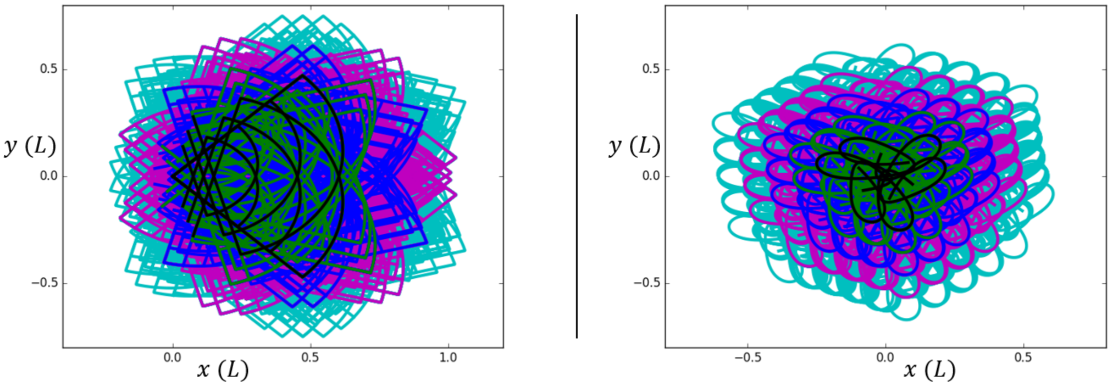

Figure 3.4: Both systems were able to explore their local environments in a way that is unrestricted to translation in the plane. We plotted paths to show the number of steps required to arrive at a target pose, projecting out the orientation (θ) component of the full $\mathsf{SE}(2)$ pose. At 5 steps (cyan), the system had a strong variety of poses at its disposal. We plotted motions available in 5 steps (1=black, 2=green, 3=blue, 4=magenta, and 5=cyan) Both systems appear to be capable of navigating through environments with sparse obstacles.

of [HC10; HC13; RH16].

3.4.1.1 Two-slider swimmer model:

The two-slider swimmer in Figure 3.3 moves via the prismatic joints driven by strictly positive displacements r_1 and r_2. The viscous force on each sphere is linear in translational velocity and cubic in rotational velocity. The full model is:

$$\begin{bmatrix} 3d & 0 & 0 \\ 0 & 3d & 0 \\ -dr_2 & -dr_1 & d({r_1}^2 + {r_2}^2) + \frac{d^3}{4} \end{bmatrix} \mathring{g} = R(\alpha) \begin{bmatrix} 0 & d \\ -d & 0 \\ 0 & 0 \end{bmatrix} \dot{r} \qquad (3.3)$$

$$\alpha = -\frac{\pi}{2}$$

where R takes input parameters to a rotation about the origin on $\mathsf{SE}(2)$.

3.4.1.2 Three-branch swimmer model:

We also designed the three-branch swimmer (see Figure 3.3), another viscous swimmer. Three links are free to rotate from the points of the triangle. For biological intuition for how a system like this might move, a starfish might move like a pentagonal five-branch system with longer segments of links at each vertex. The links interact via the slender body theory of Cox [Cox70], the same that was used for the swimmer in [HC13] and paddles in [KBR19]. The drag of the triangular piece is represented by three static links that point from the center of the triangle to their respective attachment points.

3.4.2 Hand selecting gaits

3.4.2.1 Gait selection for two-slider swimmer:

By inspection of the connection vector field of the rotational component of the two-slider swimmer (see Figure 3.3), we saw that a variety of turning modes could be excited. Hand-selected gaits all started at the origin of the base space, travel along the axis of one shape variable, then translated at a constant radius from the origin, traveling from one positive end of a shape axis to the other. Each path was then sent to the origin via the other shape variable. We can see from the curl of the vector field that clockwise gaits will yield positive rotation, and counter-clockwise gaits will yield negative rotation. The three paths printed in red represent three magnitudes of turning the system can choose. The larger the radius, the greater the turn will be, as explained in Figure 3.3. Each gait also induces a translational displacement of the system from its starting location.

3.4.2.2 Gait selection for three-branch swimmer:

The three-branch swimmer is less amenable to inspection by the connection vector field methods since it has a third shape variable. Reduction methods (such as [Rie+19]) can make such gait analysis useful for more complex Stokesian systems. Geometric gait optimization can also be employed on this analytical system to obtain a collection of gaits, maximizing various objective functions [RH19]. We avoided these methods to reinforce that principled gait design concepts are not needed to design a functional gait library for the swimmer. Here, we selected three gaits primarily for their symmetrical features, enclosed shape space volume, and aversion of self-intersections. Two links oscillated in anti-phase, providing a thrust that acts through a line from the midsection of their attachment points to the third link's attachment point. The third link oscillated out of phase by a quarter cycle. We designed the gaits as:

$$r_{mod(k,3)+1} = \sin(\varphi) \tag{3.4}$$

$$r_{mod(k+1,3)+1} = 1 - \cos(\varphi) \tag{3.5}$$

$$r_{mod(k+2,3)+1} = -1 + \cos(\varphi) \tag{3.6}$$

for $\varphi \in S^1$ with gaits γ_k enumerated k = (1,2,3). These three gaits generate three group actions, which can also be run backward in φ, generating three inverse group actions.

Each system had six gaits at its disposal. By inspection of Figure 3.4, we observe the local planning ability of the systems, only using the six gaits as possible actions (letters) of their total motion (word). We highlight the key takeaway of this section. Behaviors that were not useful in isolation were critical to providing dense coverage. Furthermore, these behaviors may lie outside the scope of typical behaviors that a roboticist may prescribe for a system.

3.5 Connecting Gait and Motion

The algebraic structure for computing available motions is straightforward. Separate gaits were concatenated as a string of group multiplications. What dynamical properties were required for such assumptions? We cover the assumptions we made in this section, using the language of geometric mechanics. For general dynamical systems, combining gaits would require a transition behavior that matches the internal state $(r, \dot{r}, p)$ of the endpoint of one gait and connects it with the internal state $(r, \dot{r}, p)$ of the starting point of the next gait. There exist a class of systems where the matching requirements are highly relaxed.

3.5.1 Planning simplifications in principally kinematic systems

The class of systems we focused on in this work inhabit the Stokes regime [KM95], which encompass the dynamic qualities of the principally kinematic case covered in [OB98]. A well known example of such systems is low Reynolds number swimmers [KM]. However, we recently accumulated evidence that this theory applies to multi-legged locomotion [WZR19; CHG20]. The function $A(\cdot)$ connects gaits, as body shape loops, to the motion they induce, called the “holonomy” of the loop.

It’s known from Stokes Theorem that a closed loop integral on a vector field is equal to the area integral of the volume enclosed by the loop. This theorem extends to higher dimensional spaces and provides the flexibility of inducing equivalent group actions no matter where the closed loop starts or stops. Furthermore, any path in the kernel distribution of $A(\cdot)$ can connect such loops to one-another without introducing an additional motion in the group. In practice, however, obtaining this kernel from data requires cumbersome sampling and system identification.

3.5.2 Representational simplifications

Given a gait γ_b, the body frame motion M_b it produces could, in principle, be a function of the initial point in the gait cycle and the speed with which this cycle is executed. For systems where momentum is dominated by friction or constraints (Stokesian systems), this is not the case. In those systems there exists a map $A(r) : r \in R \mapsto \mathsf{L}(\mathsf{T_r}, \mathsf{g})$ taking shapes to linear maps from shape velocities to the Lie Algebra g of G. This leads to the "reconstruction equation" $\dot{g} = L_g A(r) \dot{s}$ where $L_g : \mathsf{g} = \mathsf{T}_e \to \mathsf{T}_g$ is the lifted left action of the group element g. Thus, if two base loops are connected at any point, the combination of their actions can be represented as a group multiplication of their respective g elements.

There are infinite ways to take a gait library and coordinate it into a complete motion planner. Typically people have a scheme for transitioning between gaits. The overhead of finding such transitions for systems with no model can be large.

When selecting a collection of gaits for computing coverage, we required that each shares a point in the base space and thereby allowing gait cycles to be applied in any order.

3.6 Setup for Discovering a High Coverage Gait Library

To illustrate the approach on a more classical system, we simultaneously optimized three gaits on Purcell swimmers to provide coverage of a portion of $\mathsf{SE}(2)$ surrounding the identity using their $h_4(\cdot)$ cost.

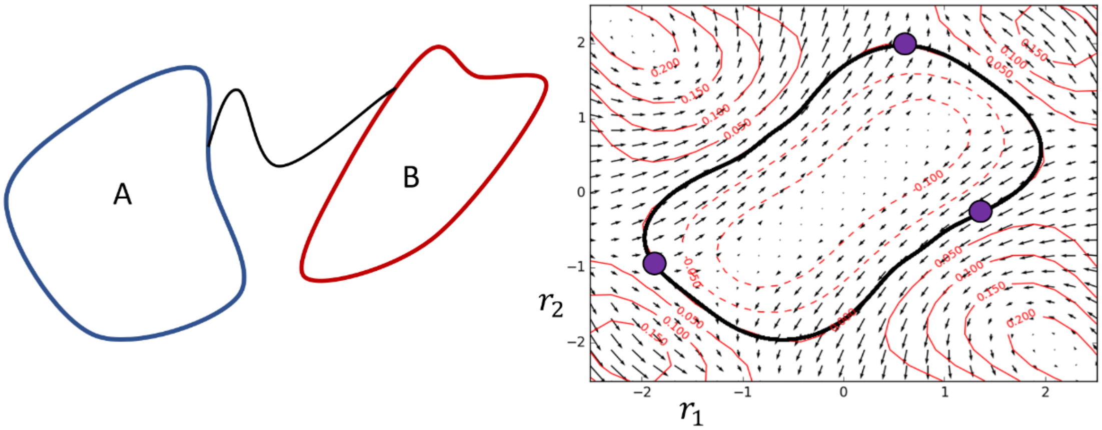

Figure 3.5: When composing motions, one has to consider the sequencing of gaits (like A and B pictured left) that may be separated in the shape space R. Planning in the Stokes regime offers some convenient structure for the composition of motions. In this regime, cycles in the internal state generate group motions irrespective of the point along the cycle that the motion starts. For example, the extremal gait (bold black line on the right) for the three-link Purcell swimmer can be started at any point on the loop (such as the purple markers). Execution of a cycle from any point will generate the same body motion. This structure greatly simplifies requirements for sequencing motions on principally kinematic systems.

3.6.1 Coverage point selection

This coverage point distribution included equally weighted points derived from all possible combinations of the following values, totaling 125 points:

$$x = [-1, -0.5, 0, 0.5, 1] \tag{3.7}$$

$$y = [-1, -0.5, 0, 0.5, 1] \tag{3.8}$$

$$\theta = [-\pi, -\frac{\pi}{2}, 0, \frac{\pi}{2}, \pi] \tag{3.9}$$

where units for translation were body lengths and units for rotation were in radians. These spanned the translational bounds of moving by one body length and the rotation bounds of rotating by a half of a full rotation.

3.6.2 Model extraction and motion parametrization

A single iteration of learning involved experimentally running each of the three gaits for 30 noisy cycles, modeling and optimizing their motion via the framework of [BHR18]. We parametrized the gaits in a modified version of the ellipse with bump function parametrization also used in [BHR18]. The following parametrization p is a modification that allows the base point, b_i, of the three gaits to be an explicit parameter:

$$\begin{aligned} r_i(t) :=& c_i + (b_i - c_i)\cos(\Omega t) + a_i \sin(\Omega t) + \\ & \sum_{k=0}^{N_o-2} u_{i,k}\, w\left(t - k\frac{2\pi}{N_o}\right) \end{aligned} \tag{3.10}$$

$$w(x) := \begin{cases} 1 + \cos(xfN_o) & |xfN_o| < \pi \\ 0 & |xfN_o| \geq \pi \end{cases}, \tag{3.11}$$

with gait parameters

$$p_i = (c_i, b_i, a_i, u_{i,k}). \tag{3.12}$$

In this work, we used $N_o = 18$ and $f = 3$, totaling 16 bumps. Two bumps were elided ($k = 17, 18$) via this representation such that base point b is left unshifted.

3.7 Finding Coverage with Purcell swimmers

Our first investigation was to see how well Purcell swimmers can optimize three gaits simultaneously for the uniformly distributed set of coverage points of §3.6.1. We observed how the ability to optimize these gaits changed as we added joints to the swimmer. We started with two joints (the three-link Purcell swimmer) and built our way up to eight joints. We repeated the optimization process 30 times for each swimmer.

At the beginning of each optimization, a random joint was stimulated with a sine wave. The stimulated joint was distinct for each gait. The only exception to this was that for the three-link Purcell swimmer, a gait had to be repeated since there were three initial gaits and only two joints. The swimmers used 30 cycles at each gait to build a model. Then, the swimmers used the models provided by [BHR18] to predict how changing the parametrization of their three gaits could be combined to optimize a 4-step plan over the coverage points provided[1]. An iteration of the optimization involved stepping along the policy gradient of three gaits (step size computed via [BHR18]) and simultaneously updating the three gaits. The results are recorded in Figure 3.6.

We ended the optimization after 30 iterations. The test showed that the swimmers were able to use the coverage metric to consistently find a gait library for local motion planning. Having two joints were sufficient to find a functional library, but having

---

[1] Using 4 steps allowed us to include knowledge of the commutator motions noted in §3.3.1.

three joints presented a notable improvement. This jump in performance was less surprising after considering that the third joint allowed the swimmer to become fully actuated (when in non-singular configurations) with respect to $\mathsf{SE}(2)$. After the third joint was added, the convergent behavior of the swimmers was consistently within the performance noise window of adding another joint.

The convergence rate of the swimmers improved when adding the third and fourth joints. For all swimmers containing 3 or more joints, the standard deviation of performance reached $h = 0.4$ by the tenth trial. Here, we calculated h as the average normed distance to a coverage point. Converging at this expedient rate required exactly 900 cycles of robot data. If we ran the robots at 3Hz, the optimizations would have converged after collecting just five minutes of experimental data, even on the nine-link swimmers.

3.8 Investigating the Ability of the Purcell Swimmer to Recover from Joint Locking

In trials 30-60 of Figure 3.6, we tested the ability of the Purcell swimmer to recover from “injury”. We took the optimal collection of gaits from the first 30 trials and found the joint that used the highest amplitude behavior. We locked this joint at its value taken at the base point of the parametrization. We then asked the swimmers to take this optimal policy from the first 30 trials and use it to find a new optimal policy despite the fact that one of its joints had been turned into a static kink.

The three-link swimmer was unable to move as a result of the injury, a consequence of the scallop theorem in Stokesian systems [Lau11]. The four-link swimmer was able to partially recover. It was equipped with two functional joints, yet was unable to achieve the coverage scores of the un-injured three-link Purcell swimmer. This may suggest that the injury resulted in a body geometry that was less equipped to

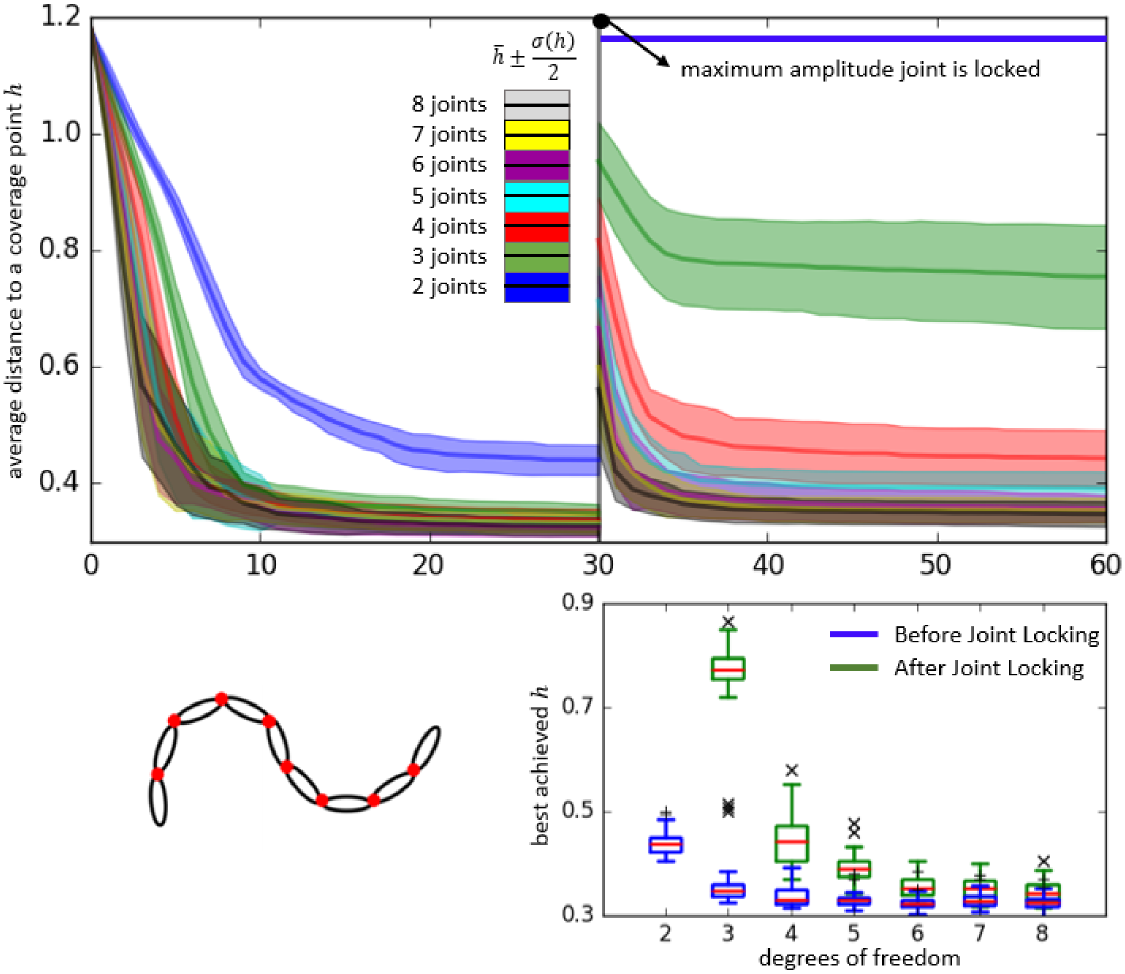

Figure 3.6: Purcell swimmers of varying complexity, such as the nine-link (pictured bottom left) were optimized for three gaits that maximize coverage. See §3.6.1 and §3.6.2 for details on the setup of the experiment. We plotted the mean (top, solid lines) and standard deviation (transparent bands) over 30 separate simulations of the average distance of goal motions to the nearest available motion, denoted h. We can see how h changes across trials and the number of joints used by the swimmer (2=blue, 3=green, 4=red, 5=cyan, 6=magenta, 7=yellow, 8=black). At iteration 30 (marked by a vertical grey line), we plotted how well the swimmers adapt to having the maximal amplitude joint locked. We also observed how the quality of the coverage of the library varies by the number of joints used by the swimmer (bottom right) before (blue box plots) and after (green box plots) joint locking.

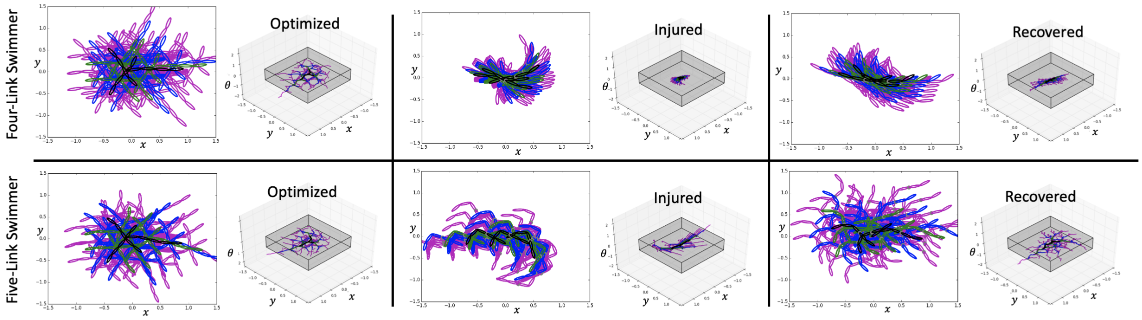

Figure 3.7: This provides a detailed look at two optimization process for a four-link and five-link swimmer in the study summarized in Figure 3.6. We plotted the 4 step horizon (1=black, 2=green, 3=blue, 4=magenta) at various trials on the plane (left in each section) and on $\mathsf{SE}(2)$ (right in each section). For reference, we plotted the unit volume in $\mathsf{SE}(2)$ (gray box) over which the coverage points were uniformly distributed. For the four-link swimmer, we showed the optimal policy before injury in trial 22 (top right), the consequence of a locked joint (grey dot) on the optimal policy in trial 30 (top middle), and the optimal policy recovered while the joint remains locked in trial 52 (bottom right). The four-link swimmer was strongly impeded in its ability to recover a high coverage collection of gaits post-injury. For the five-link swimmer, we showed the optimal policy before injury in trial 17 (top right), the consequence of a locked joint (grey dot) on the optimal policy in trial 30 (top middle), and the optimal policy recovered while the joint remains locked in trial 54 (bottom right). The five-link swimmer was not impeded in its ability to recover a high coverage collection of gaits post-injury.

optimize coverage than the conventional three-link Purcell swimmer. As more joints were added, the redundancy of joints both minimized the dynamical impact of injury and provided a larger space of solutions for recovery.

One might pose the question: How many joints does the Purcell swimmer need to be robust to a single joint locking? One could make a case that four joints were a justified selection. The five-link swimmers were notably better at finding high coverage libraries during recovery than the four-link swimmers and remained within the standard deviation of performance of the six-link swimmers. The top row of Figure 3.7 details one optimization process for a swimmer with three joints. It is clear that before injury, the swimmer was able to achieve local poses. The injury greatly handicapped this ability, even with the opportunity to recover. Likewise, the bottom row of Figure 3.7 details one optimization process for the swimmer with four joints. Before injury, the swimmer also found a useful gait library. The injury clearly hindered its ability to move, but given the opportunity to recover, the swimmer persisted in finding a new collection of behaviors that adapt to cover the local space.

As a result of this analysis, one could make the claim that up to adding a fourth joint, redundancy only aided the optimization process. This is a rare statement to make since adding a degree of freedom typically involves a substantial increase in sampling requirements, both lengthening the convergence rate and making it less certain. In this example, convergence rate was either improved or stays approximately the same as joints are added. Here, combining the methods of [BHR18] and the coverage metric allowed redundancy in the internal state to be an asset for behavior optimization rather than a liability.

3.9 Implementation on Hardware

Here we communicate the general and noise-robust qualities of our approach by optimizing coverage on real hardware with an unknown model. We did not have

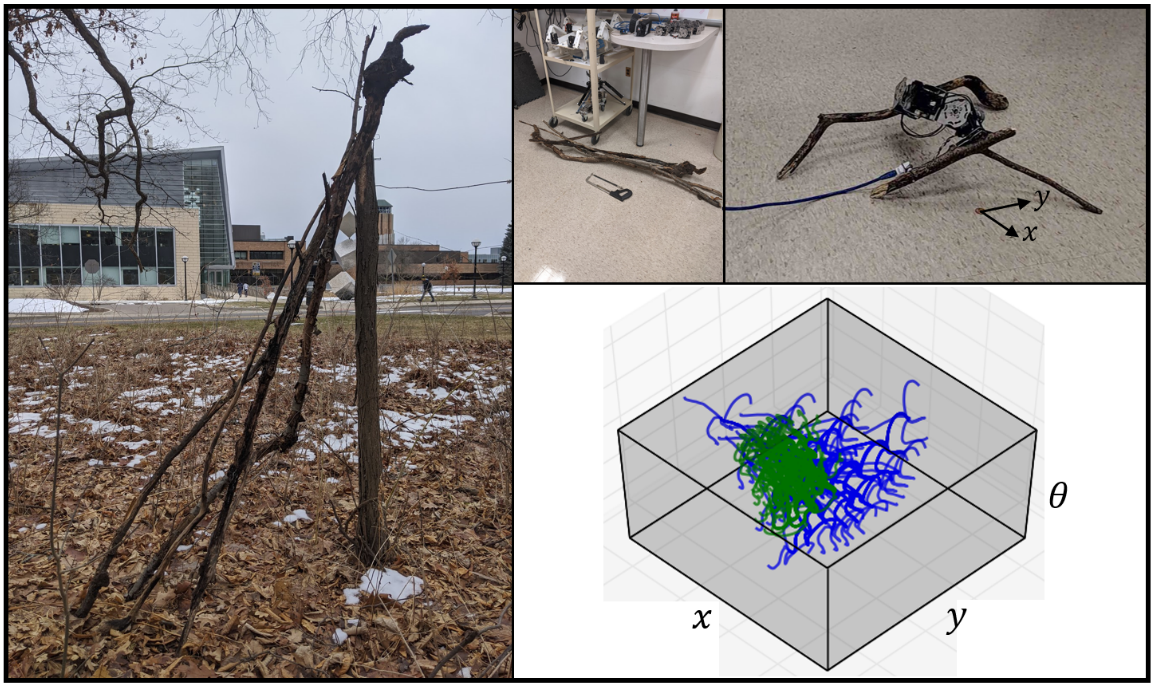

Figure 3.8: This robot (top right) was built from dynamixel modules and tree branches available nearby (left and middle left). The trajectories showcase the available 1 to 4 cycle motions of the system (bottom right) from the robot's origin before (green) and after (blue) the coverage optimization.

explicit knowledge of the kinematics, mass distribution, or material properties of this system when running the modeling and optimization algorithms.

3.9.1 Methods on hardware

Day 1: Inspired by the hardware used in [Mae+18], we traveled into the woods foraging for tree branches. We gathered such elements, transferred them to the lab, and sectioned the branches into viable robot appendages. We then constructed a robot by fixing tree branches to the endpoints of a chain of three Dynamixel actuator modules.

Day 2: We equipped the robot with motion capture markers and connected it to a computer running the gait modeling and optimization algorithms. We used the same equipment that was used in 2.9.1. To build a physics model centered at a given gait, we collect 20 cycles of noisy input data on the robot and fit a regression informed by physics and geometry [BHR18]. We then compared the outcomes of two different optimizations for the (x, y, θ) outcomes of a gait or gaits, taking the position of the robot prior to application of a gait cycle as the origin.

(1) Find one gait to move forward without turning: We designed a gait optimization to maximize $x - y^2 - \theta^2$ (per cycle) given the coordinates of Figure 3.8 and units of body lengths ($\frac{1}{3}m$) and radians.

(2) Find three gaits that optimize coverage in a volume of $\mathsf{SE}(2)$: Given the ability to use the 3 gaits in up to 4 combined cycles, we designed a gait library optimization to minimize the distance (computed on the Lie group), from 125 points distributed across all combinations of coordinates $x = [-1, -\frac{1}{2}, 0, \frac{1}{2}, 1]$, $y = [-1, -\frac{1}{2}, 0, \frac{1}{2}, 1]$, $\theta = [-\frac{\pi}{2}, -\frac{\pi}{4}, 0, \frac{\pi}{4}, \frac{\pi}{2}]$. The gray volume in Figure 3.8 contains all of the coverage points.

3.9.2 Results on hardware

For the first goal function, we seeded a zero motion gait oscillating the middle joint with a sinusoidal input. We executed 15 iterations of our data-driven gait optimization algorithm, each consisting of 20 cycles of motion. Running at $\frac{1}{2}$Hz, each trial took 40 seconds. After the 8th iteration, the robot was able to travel 40% of its body length per cycle with a turning rate of 0.10 radians per cycle.

For coverage, we first completed an exploratory sampling of motions (12 cycles). From these 12 different gaits, we selected the subset of 3, which performed the best on the coverage metric. After 5 iterations of trials (60 cycles, 20 for each gait), the system found a more complementary set of gaits improving the coverage score from 0.97 to 0.76 average distance to the coverage points.

3.10 Discussion and Conclusions

In this paper, we introduced a new metric for the optimization of robot motions. This metric involved calculation of the composition of motions from a small library of primitives, determining their utility in "covering" some region of the local position space, formulated as a Lie group. What is novel about this approach is that

- It eliminated human bias from prescribing a limited set of allowable primitives for a robot.
- It allowed for the use of unconventional robot designs for navigation.
- It allowed malfunctioning robots to quickly recover the ability to move through space.

We showed the Purcell swimmers' ability to recover from injury using the data-driven geometric gait optimizer, guided by the coverage metric. Some interesting trends emerged during these tests. The swimmers converged to a high coverage gait library (containing three gaits) despite variation in the number of links and initial gaits. This suggests insensitivity in the gait optimization when using the coverage metric.

Furthermore, coverage allowed us to investigate the role that redundancy might play in the ability of the swimmers to recover high coverage gait libraries post-injury. We found that around four degrees of freedom, the addition of a joint no longer provides a substantial change in the ability of the swimmer to recover. The ability to apply this analysis to other robots could help inform what degree of complexity is appropriate when designing a robot.

Finally, the coverage optimization on a robot made of tree branches was unable to find gaits that translate without substantial rotation. The tree branch robot was able to find a useful portfolio of maneuvers for navigation on a timescale that is competitive with an implementation of reinforcement learning by Google [Ha+20].

The tree branch robot example speaks to the morphology agnostic properties of data-driven geometric gait optimization. The tree branch robot used in this framework could be substituted with robots of many unexpected forms. As long as the system acts near the Stokesian regime of locomotion, the methods of Chapter II assist in building behavioral models that inform performance improvements.

CHAPTER IV

Data-Driven Geometric Modeling Extensions

4.1 Modeling Extensions to Shape-Underactuated Dissipative Systems

4.1.1 Motivation

Rigid, fully actuated mechanisms are emblematic of the classical field of robotics. The development of passive elements [PW; RMH14; KRC15; AB97; SBK01] and soft actuators [KCH99; Tol+14; Seo+13; Pel00; SBC15] offers the potential for breakthrough improvements for the design of future systems. Passive elements have the potential to assist in designing mechanisms that are safer, cheaper, more energy efficient, and more resilient to impact damage. However, these design improvements typically come at the cost of precise control of the internal state of the system. The degree of underactuation of internal state and the complexity of soft mechanisms can both exacerbate this problem.

Early robotics research showed that a convenient way to add compliance to a mechanism is to add a spring in series with a motorized joint [PW]. The "Series Elastic Actuator (SEA)" was introduced to humanoids [Rad+15] and snake robots [Rol+14] with the goals of providing compliant, torque controlled interaction with the environment and higher damage resilience. The design advantages of SEAs come at the expense of high-bandwidth position control. It became difficult to execute precise body-shape trajectories that would be possible in the fully actuated, otherwise identical, systems. In robots with soft actuators, the shortcomings in position control are exacerbated by the sensitive nonlinear models of pneumatic devices, dielectric elastomers, and other soft actuation techniques [WJ10; RT15]. The challenges of precise fabrication and assembly make it difficult to reliably reproduce dynamical outputs across copies of these devices.

The elusive nature of obtaining predictive models for highly underactuated systems shares features of the challenges of modeling their high degree of freedom, fully actuated counterparts. For fully actuated dissipative systems, we have previously published sample-efficient techniques to model locomotion systems with noisy shape control using cyclic behavioral data [BHR18; KBR19]. Seminal work by Shapere, Wilczek, Marsden, Kelly, Ostrowski, Bloch and others [SW89b; KM95; Blo+96; MO98; OB98; CMR01] showed that the Newtonian physics of locomotion can be refactored into a kinematic term (the mechanical connection of [MO98]) and a momentum term. At the limit of large friction, the momentum term disappears, leaving a class of models which we have shown to be easy to system identify [BHR18]. Further, with finite-but-large dissipation, the influence of momentum can be folded into a nonlinear correction to the connection, with only a small increase in the complexity of the model identification process [KBR19]. Thus models for predicting the influence of shape input on body velocity can be built strictly from observation without any mechanical analysis specific to the system – all that is needed is “sufficiently rapid” dissipation of momentum.

In the current work, we extend these ideas to underactuated systems. First, we identify the class of “Shape-Underactuated Dissipative Systems” (see §4.1.3) to which our methods apply. Informally, these are systems that have fewer actuators than internal degrees of freedom and whose mechanics are governed primarily by frictional and damping forces, rather than inertial ones. We claim that SUDS are a highly useful and broad class of dynamical systems in practice. We then show how data-driven geometric modeling techniques can be extended and used to identify predictive models for SUDS (see §4.1.3). For the subclass of SUDS whose internal dissipation is linear, the technique further allows us to collapse our model complexity, achieving a complexity that grows linearly in the degree of underactuation (see §4.1.4). To demonstrate the efficacy of our approach, we examine its performance on simulated viscous swimming data (see §4.1.5), validating that predictive SUDS models can be identified for soft, high dimensional systems with small amounts of trial data. Finally, we discuss the relevance of SUDS identification in modern robotics applications.

4.1.2 Background: Data-Driven Connection Modeling

In the field of geometric mechanics, the equations of motion arise from dynamical constraints derived from Lagrangian or Hamiltonian descriptions, after which group symmetries are applied to generate a reduced form [CMR01; Blo+96]. The representation of these equations incorporates the uniformity of the operating environment. This involves a systematic reduction of the dynamics, achieved by quotienting the dynamics by its dependence on group. A common and representative case is the symmetry expressing the fact that a body's interactions with a uniform environment do not depend on its position and orientation in that environment[1]. Under these circumstances we can re-write the equations of motion using a "reconstruction equation"[OB98]. This appears as

$$\begin{aligned}\mathring{g} &= A(r)\dot{r} + \mathbb{I}^{-1}(r)p \\ \dot{p} &= f(r, \dot{r}, p). \end{aligned} \tag{4.1}$$

---

[1] While our work applies without modification to other Lie group symmetries, we will tacitly assume that the symmetry is a subgroup of $\mathsf{SE}(3)$ and use the terms "body frame" and "body shape" for the "fibre" and "base space projection" that appear in the fibre bundle formulation of this theory.

These tools express in a formal and complete way the intuition that symmetry in the environment should allow us to write equations of motion relative to the body frame.

As was shown for the case of the reduced Lagrangian, one can separate the influence on body frame motions into two factors, a kinematic contribution and a momentum contribution. Particularly, when one of these contributions dominates the other, we gain strong insight into the key influences and features of the locomotion model. They also introduce a significant simplification – the momentum p appearing in them is of dimension equal to that of the group. In the general case, this reduces the number of dynamical equations by the dimension of the group, since $\overset{\circ}{g}$ is now an output rather than a state. More profoundly, since in the realm of robotics the body shape $r(t)$ can often be dictated with high-gain feedback, the dimension of the remaining equations is the dimension of p.

When the motion is governed by linear constraints on the velocity, the dimension of p further reduces; these are sometimes known as "Pfaffian constraints". For moving systems with environmental symmetries, Pfaffian constraints often come in the form of body frame velocity constraints (e.g., no sideways slipping). Friction, in the form of a Rayleigh dissipation function, can further dissipate the momentum $p \to 0$, and if it does so quickly enough, the results are similar to those of a system governed by Pfaffian constraints. With momentum gone, the equation retains only the $A(r)\dot{r}$ term, known as the "mechanical connection" [MO98]. These systems are "principally kinematic" in the sense that their motion depends only on the shape of their body configuration curve, but not on the rate.

The most well known, principally kinematic locomotors are viscous swimmers acting in low Reynolds environments [HC13]. By exploiting the structure of the mechanical connection, tools have been developed for coordinate system selection, gait identification, and behavioral optimization [HC11a; WL13; HC15; WO16; RH19].

Predictive global models are often challenging to obtain for real animals and for physical hardware. System identification techniques [HC13; Dai+; Sch+19; Ast+20] allow for data-driven modeling of animals and robots but require a large amount of experimental data. Typically some reduction of the representation of the shape space is needed to make these methods produce tractable models of complex animals and robots. Thus, there is a real need for modeling techniques with lean data requirements that can handle high dimensional representations of the body shape.

In [BHR18], we developed a data-driven approach to geometric modeling and optimization. It allows us to identify a mechanical connection that governs a rhythmic motion with very little data (e.g. on the order of 30 cycles for a nine-link Purcell swimmer). We built this estimation framework by combining oscillator theory [Rev09; RG08; RK15b] and geometric gait optimization [RH17; RH19]. Using a phase estimator from [RG08], we computed phase from observed cyclic shape data. Grouping measurements by phase allowed us to compute a Taylor series approximation of the mechanical connection at each phase using linear regression. Further theoretical analysis showed that when momentum decays quickly but not instantly, there exists a nonlinear $A(r, \dot{r})$ close to the linear mechanical connection; this additional nonlinearity was straightforward to capture with the inclusion of additional terms of the order of the momentum decay time-constant [KBR19].

4.1.3 Shape-Underactuated Dissipative Systems (SUDS)

The locomotion model for systems whose dynamics have the structure of a mechanical connection take the form:

$$\mathring{g} = A(r)\dot{r} \tag{4.2}$$

where $r \in \mathbb{R}^n$ spans the shape space R, g is an element of a Lie group G, and $A(r)$ is an infinitesimal lift from shape velocities to body velocities. The notation $\mathring{g}$ denotes the world velocity $\dot{g}$ written in the body frame, computed as $g^{-1}\dot{g}$ for matrix Lie groups.

Previous work [HC13; HC15] showed that for mechanical connections dominated by drag, the internal wrenches along the degrees of freedom of the shape can be written as:

$$\tau = -\mathcal{M}(r)\dot{r}, \tag{4.3}$$

where $\mathcal{M}$ is a Riemannian metric of the shape space that weights the cost of changing shape in various directions. Because $\mathcal{M}$ is positive definite, its negation in equation 4.3 means that the system is "passive" in the sense used in control theory – changing shape always consumes energy.

For underactuated systems, arbitrary choice of instantaneous shape velocity $\dot{r}$ is infeasible. Consequently, the form of equation 4.2 is not directly useful for planning system motions. We split the shape configuration and force vectors as

$$r = r_u \oplus r_p \qquad \tau = \tau_u \oplus \tau_p \tag{4.4}$$

where u indicates controlled degrees of freedom and p indicates passive degrees of freedom. These passive degrees of freedom are governed by some dynamical relationship in which the wrench on the passive joint is a function of

$$\tau_p = \tilde{f}(r, \dot{r}, \mathring{g}). \tag{4.5}$$

We substitute equation 4.2 into equation 4.5 to reduce this relationship to a mapping from shape and shape velocity to the internal wrenches

$$\tau_p = f(r, \dot{r}). \tag{4.6}$$

The $u \oplus p$ splittings of r and τ break $\mathcal{M}$ into four blocks

$$\mathcal{M} = \begin{bmatrix} \mathcal{M}_{uu} & \mathcal{M}_{up} \\ \mathcal{M}_{pu} & \mathcal{M}_{pp} \end{bmatrix} \tag{4.7}$$

where for brevity we supress the dependence of $\mathcal{M}$ on r. We now can represent the passive wrenches in two ways, drawing from equations 4.3 and 4.6, such that

$$\tau_p = -\mathcal{M}_{pu}\dot{r}_u - \mathcal{M}_{pp}\dot{r}_p = f(r, \dot{r}), \tag{4.8}$$

and after rearranging,

$$-\mathcal{M}_{pp}\dot{r}_p = f(r, \dot{r}) + \mathcal{M}_{pu}\dot{r}_u. \tag{4.9}$$

Noting that many physical systems of consequence exhibit linear or nearly linear dissipation, we add the assumption that we may rewrite f as an r dependent affine function of $\dot{r}$,

$$f(r, \dot{r}) = f_o(r) + F(r)\dot{r} = f_o + F_u\dot{r}_u + F_p\dot{r}_p. \tag{4.10}$$

Combined with equation 4.9, we arrive at an equation where each term is constant or linear in shape velocity

$$-\mathcal{M}_{pp}\dot{r}_p = f_o + F_u\dot{r}_u + F_p\dot{r}_p + \mathcal{M}_{pu}\dot{r}_u. \tag{4.11}$$

This expression is equivalent to

$$-(\mathcal{M}_{pp} + F_p)\dot{r}_p = f_o + (F_u + \mathcal{M}_{pu})\dot{r}_u, \tag{4.12}$$

which allows us to show that $\dot{r}_p$ can be written in a form that is affine in $\dot{r}_u$.

Now we show that $(\mathcal{M}_{pp} + F_p)$ is full rank, which will prove that the affine relationship between $\dot{r}_p$ and $\dot{r}_u$ is not degenerate. Term $\mathcal{M}_{pp}$ is positive definite since it is a diagonal block of $\mathcal{M}$, which we have established is itself positive definite. Term F_p is semi-positive definite since any damped system will have a non-negative power dissipation from damping $\dot{r}_p^T F_p \dot{r}_p$. The sum of a positive definite matrix and a semi-positive definite matrix is itself positive definite, and thus $(\mathcal{M}_{pp} + F_p)$ is invertible.

Because equation 4.2 is linear (and thus affine) in $\dot{r}$, and $\dot{r}_p$ is affine in $\dot{r}_u$, we obtain that $\overset{\circ}{g}$ must be affine in $\dot{r}_u$. The equations for $(\overset{\circ}{g}, \dot{r}_p)$ are affine in $\dot{r}_u$:

$$\overset{\circ}{g} = A_u(r)\dot{r}_u + \overset{\circ}{g}_o(r) \tag{4.13}$$

$$\dot{r}_p = -(\mathcal{M}_{pp} + F_p)^{-1}\left[f_o + (F_u + \mathcal{M}_{pu})\dot{r}_u\right] \tag{4.14}$$

In many control applications the control input is τ_u rather than $\dot{r}_u$. Using equation 4.3 we can solve by substituting equation 4.14 to give an explicit affine formula for τ_u from $\dot{r}_u$

$$\tau_u = -\mathcal{M}_{up}\dot{r}_p - \mathcal{M}_{uu}\dot{r}_u \tag{4.15}$$

We define a “Shape-Underactuated Dissipative System (SUDS)” as a mechanical system operating within the dynamical constraints of equation 4.2 and equation 4.3. We focus on SUDS containing linear passive elements of the constrained form given by equation 4.10. These systems are therefore governed by motion models comprised of equations 4.13 and 4.14. When combined these equations lead to the observation that

$$(\overset{\circ}{g}, \dot{r}_p)^T = \tilde{C}(r) + B(r)\dot{r}_u, \tag{4.16}$$

e.g. the dynamics of SUDS are a nonlinear function of shape r, affine in the directly controlled shape velocity $\dot{r}_u$.

4.1.4 Estimation for SUDS

Now that we have established a dynamical characterization of SUDS, we can discuss the ramifications of this characterization for the estimation of motion models from data. If analytical models are available, methods derived in [RH20] provide a way to perform gait optimization on drag dominated systems with an elastic joint. However, when analytical models are not available, sample efficient methods for system identification are required for data-driven behavioral optimization. We will show that the characterization presented in §4.1.3 will be important for data-efficient system identification of highly underactuated systems. Following the approach we took in previous work [BHR18], we focus on identifying the dynamics within a "tube" around a nominal trajectory θ by expressing the shape as $r := \theta + \delta$. Here δ expresses deviation from the nominal trajectory. We then consider the approximation of $(\overset{\circ}{g}, \dot{r}_p)$ by a first-order Taylor expansion in $(\delta, \dot{\delta})$ as

$$\begin{aligned}(\overset{\circ}{g}, \dot{r}_p) \approx &\tilde{C}(\theta) + \frac{\partial \tilde{C}}{\partial r}(\theta)\delta + B(\theta)(\dot{\theta}_u + \dot{\delta}_u) \\ &+ \frac{\partial B}{\partial r}(\theta)\delta(\dot{\theta}_u + \dot{\delta}_u).\end{aligned} \tag{4.17}$$

However, because $\dot{\theta}$ is a predetermined function of θ, we can combine terms (suppressing the (θ) for readability)

$$C := \tilde{C} + B\dot{\theta}_u \tag{4.18}$$

$$C_r := \frac{\partial \tilde{C}}{\partial r} + \frac{\partial B}{\partial r}\dot{\theta}_u \tag{4.19}$$

which provide the following linear regression problem at each θ,

$$(\overset{\circ}{g}, \dot{r}_p) \sim C + C_r\delta + B\dot{\delta}_u + B_r\delta\dot{\delta}_u. \tag{4.20}$$

The regression in equation 4.20 expresses the instantaneous body and shape velocities given the current shape (referenced from r, δ) and the control input (referenced by $\dot{\delta}_u$) to the system.

4.1.4.1 SUDS balance compactness of model with capability to approximate dynamics

A primary challenge of system identification is to select the unknown parameters to solve for the model governing the system dynamics. Choosing too few parameters can cause underfitting while choosing too many parameters can often cause overfitting. Here we show that the characterization of SUDS dynamics allows for a compact yet descriptive set of parameters to seed system identification. In particular, we pay attention to the ability of the parameters to remain descriptive and compact at high degrees of underactuation, which is a prevalent feature of soft systems.

The overall shape space dimension is $n := n_u + n_p$, the number of directly controlled DoF and the number of passive DoF in the system respectively. Compare now the regressors of equation 4.10 to those of a more general SUDS

1. $\delta, \dot{\delta}$ for a first-order Taylor approximation of a general SUDS, having $O(n)$ unknowns.
2. $\delta, \dot{\delta}_u, \delta \otimes \dot{\delta}_u$ for a first-order Taylor approximation of a passive Stokesian system constrained as per equation 4.10, having $O(nn_u)$ unknowns.
3. $\delta, \dot{\delta}, \delta \otimes \dot{\delta}, \delta^2, \dot{\delta}^2$ for a second order Taylor approximation of the general SUDS, having $O(n^2)$ unknowns.

Thus estimation (2) provides the structural context beyond (1) to accurately model system behavior while avoiding the $O(n^2)$ growth of estimation (3). This has a clear advantage for soft systems, which typically have a small number of control inputs and a high dimensional shape space.

4.1.5 Examples of SUDS Swimmers

To illustrate our method we examined several systems that are amenable to this estimation architecture. In these systems, a viscous ("Stokes") flow regime produced the affine constraints via Newtonian force balance.

4.1.5.1 Linear Passive Swimmer

The linear passive swimmer (first row of Figure 4.1) consists of a shape-changing "T-shaped" paddle connected to a payload volume via a spring-damper. The T shape is comprised of a horizontal bar of fixed width and length r_2, affixed to the midpoint of a vertical bar which has a fixed width and a variable height $L - r_2$. As r_2 varies between 0 and L, the change of shape of the paddle interacts with a Stokes fluid, generating reaction forces. The spring-damper connection to the payload has rest length l_k, instantaneous length r_1, spring constant k, and damping coefficient c. Due to symmetry, the linear passive swimmer exerts no torques and it is constrained to move along the x axis. The single Pfaffian constraint that drives the motion model is:

$$lc\dot{x} + cr_2(\dot{x} + \dot{r}_1 - \dot{r}_2) = 0 \tag{4.21}$$

leading to the motion model

$$\mathring{g} = \frac{-cr_2}{cl + cr_2} \begin{bmatrix} 1 & -1 \end{bmatrix} \begin{bmatrix} \dot{r}_1 \\ \dot{r}_2 \end{bmatrix} \tag{4.22}$$

(in which $\mathring{g} = \dot{x}$). The group used here $G = \mathbb{R}$ is abelian, so connection vector field (CVF) analysis provides exact solutions rather than approximations [HC11a]. This exact mechanical connection persists in the presence of shape-underactuation, which acts only to restrict what shape trajectories (and therefore group trajectories) can be expressed.

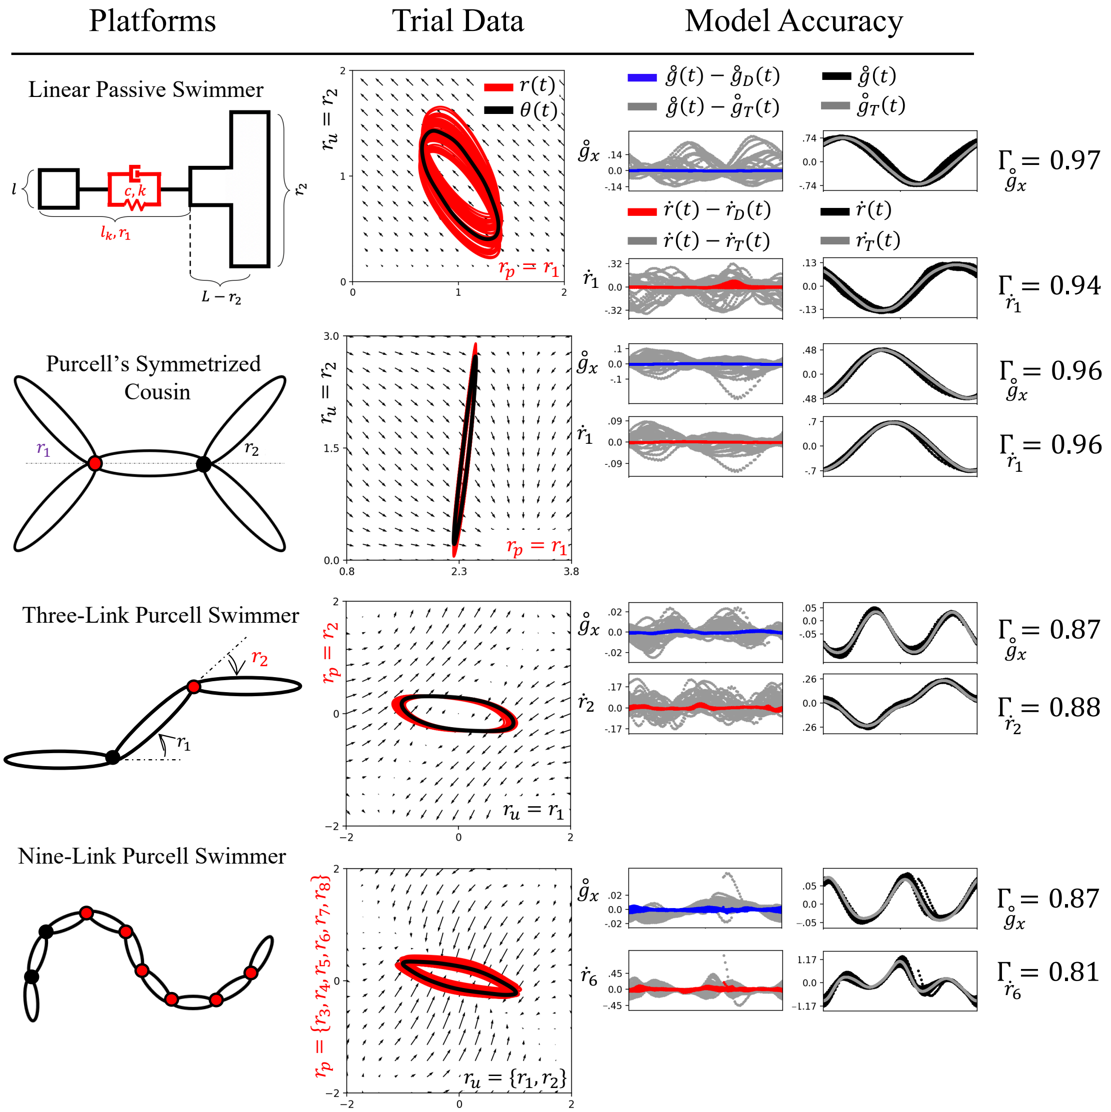

Figure 4.1: Predictive quality of data-driven SUDS models for several systems. We examined the predictive ability of regressions in equation 4.20 on simulated gait data for a linear passive swimmer, a pushmepullyou swimmer, a three-link Purcell swimmer, and a nine-link Purcell swimmer (top to bottom). In the cartoons of these systems (left column), we indicated controlled joints (black) and passive joints (red). We plotted the raw gait data (red; 30 cycles at $0.5Hz$) and the phase-averaged gait (black) for each system (second column). The metric Γ provides a reference of how accurate the data-driven connection model is with respect to the phase averaged model. We compared the two models, plotting the residuals of data-driven body velocity model (blue) and passive shape velocity (red) on top of the phase averaged model residuals (gray). We also plotted passive shape and body velocity (black) with phase averaged model indicated (yellow), demonstrating that while the phase averaged models are quite good, the data-driven connection model greatly improved the fidelity of the model, explained by the Γ metric on the right.

For this system, the internal forces can be written as

$$lc\dot{x} = k(r_1 - l_k) + d\dot{r}_1 + \omega \tag{4.23}$$

$$cr_2(\dot{x} + \dot{r}_1 - \dot{r}_2) = k(l_k - r_1) - d\dot{r}_1 + \omega, \tag{4.24}$$

where ω is the wrench that the world exerts on the system (in this case a force along the x-axis).

Combining the equations for external force balance (equation 4.22) and internal force balance (equations 4.23 and 4.24) provide three equations and three unknowns $(\mathring{g}, \dot{r}_3, \omega)$. We write the equations such that inversion of the matrix on the left-hand side will provide a locomotion model for the system's motion, given $r_1(t), r_2(t), r_3(t = 0)$. Stacking the equations, we write

$$\begin{bmatrix} cl + cr_2 & 0 & cr_2 \\ cl & -1 & -d \\ c(r_2) & -1 & (d + c(r_2)) \end{bmatrix} \begin{bmatrix} \dot{x} \\ \omega \\ \dot{r}_1 \end{bmatrix} = \begin{bmatrix} -\,c(r_2) \\ 0 \\ -c(r_2) \end{bmatrix} - \dot{r}_2 + \begin{bmatrix} 0 \\ k(r_1 - l_k) \\ -k(r_1 - l_k) \end{bmatrix} \tag{4.25}$$

The dynamics for the linear passive swimmer fit into the form of equations 4.13 and 4.14 where $r_u = r_2$ and $r_p = r_1$. As a driving signal for this swimmer, we used $r_u := 1 + \sin(t)/2$. For physical constants, we used $L = 2$, $l = 0.5$, $c = 1$, $d = 1$, and $l_k = 1$.

4.1.5.2 Pushmepullyou Swimmer

This symmetric viscous swimmer (second row of Figure 4.1), introduced in [AR08], is constrained such that the pairs of links on the left and on the right open symmetrically about the center-line of the swimmer. The symmetry allows us to assume the system moves only along the x axis. By exciting r_1 and making r_2 passive, we obtained a small forward displacement over every cycle. We chose $L = 1$, $k = 1$, and $r_k = \frac{1}{2}$.

This swimmer is also called a "pushmepullyou" swimmer, as it describes a common approach of offset motions of the left and right link pairs. The single Pfaffian constraint that drives the motion model is

$$\begin{aligned} 0 &= L\dot{x} + 2(Lc_1^2 + 2Ls_1^2)\dot{x} + 2L^2 s_1 \dot{r}_1 \\ &\quad + 2(Lc_2^2 + 2Ls_2^2)\dot{x} + 2L^2 s_2 - \dot{r}_2 \end{aligned} \tag{4.26}$$

where for brevity, denote $s_i, c_i := \sin(r_i), \cos(r_i)$ for $i = 1, 2$. This leads to the motion model

$$\dot{x} = \alpha \begin{bmatrix} -Ls_1 & Ls_2 \end{bmatrix}^T \begin{bmatrix} \dot{r}_1 \\ \dot{r}_2 \end{bmatrix} = 0 \tag{4.27}$$

$$\alpha = \frac{1}{\frac{1}{2} + c_1^2 + 2s_1^2 + c_2^2 + 2s_2^2}. \tag{4.28}$$

We place a spring on the left pair of joints such that r_1 is driven to $r_k = 0.5rad$ via spring constant $k = 1$. We write the internal torque balance on the passive joint as

$$k(r_1 - r_0) = (-2L^2\dot{r}_1 + 2Ls_1\dot{x})L + \frac{L^3}{12}\dot{r}_1. \tag{4.29}$$

This resulted in the equations

$$\begin{bmatrix} \alpha^{-1} & Ls_1 \\ \gamma_1 & \gamma_2 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{r}_1 \end{bmatrix} = \begin{bmatrix} Ls_2 \\ 0 \end{bmatrix} - \dot{r}_2 + \begin{bmatrix} 0 \\ k(r_1 - r_k) \end{bmatrix} \tag{4.30}$$

$$\gamma_1 = 2L^2 s_1 \qquad \gamma_2 = -2L^3 + \frac{L^3}{12} \tag{4.31}$$

which match the form of equations 4.13 and 4.14, where $r_u = r_2$ and $r_p = r_1$. We drove this model with $r_u := 2 + \sin(t)/2$.

4.1.5.3 Purcell Swimmer and nine-link viscous swimmer

The Purcell Swimmer and nine-link viscous swimmer (third and fourth rows of Figure 4.1) are known to have connection models [AR08]. In [BHR18], we studied the ability to model and optimize gaits with these platforms. The force balance that induces the Pfaffian constraints is presented in [HC13]. Torsional springs and dampers can act at the joints within the specified form of equation 4.10, and the model will maintain the form of equations 4.13 and 4.14. In this work, we use the model and equations of [HC13]. We use segment length $L = \frac{1}{2N}$ with a spring at each passive joint having a rest angle of 0 and a spring constant of $k_\tau = 5$. We drive the three-link Purcell swimmer with $r_u := \sin(t)$, and the nine-link Purcell swimmer with $r_u := [\sin(t), \cos(t)]$.

4.1.6 Estimator Accuracy

We sample the position and shape space of each of these systems at 100 timesteps per cycle for a 50 cycle trial. The control inputs to the system were driven by a Stratonovich stochastic differential equation, in a process identical to that used in [BHR18]. In summary, this process involves an input that is perturbed via Brownian noise while being exponentially attracted to a reference signal. The reference is periodic, defining the gait or limit cycle that the system is perturbed about. We select gaits for each system such that they noticeably excited the passive degrees of freedom. We drive each gait at a half Hz frequency since this was sufficient to produce excitation across all mechanisms. We compute each data-driven model by fitting the regressions equation 4.20 to the trial data use the same method of as [BHR18] (a fairly naive least squares regression approach).

To assess the quality of our data-driven models, we compare our SUDS regression models with the predictions obtained from a phase-averaged behavior of the same system. Such phase-averaged behaviors can be viewed as the simplest "template" model of the dynamics, whereby all periodic locomotion gaits can be viewed as oscillators [Sei+17]. We employ the *phaser* algorithm of [RG08] to reconstruct a phase from the "observation" data produced by the simulation, as this algorithm has been shown to be effective in producing phase driven models for many animal and robot locomotion systems [Mau+15; Wil+17; CR20]. In the sequel, we denote by $\mathring{g}$ and $\dot{r}$, the ground truth body velocity and shape velocity samples (respectively). By $\mathring{g}_T$ and $\dot{r}_T$, we denote the predicted value for these quantities projected onto the phase model of the system [2]. Finally, $\mathring{g}_D$ and $\dot{r}_D$ denote the data-driven model-predicted values of these same variables.

We define an accuracy metric for our predictions as one minus the ratio of the error in the data-driven prediction to the error in the phase-only predictions,

$$\Gamma_* = 1 - \frac{\sum_{i=1}^{m} \lvert *_D - * \rvert}{\sum_{i=1}^{m} \lvert *_T - * \rvert}, \tag{4.32}$$

for m samples and $* = \{\mathring{g}, \dot{r}\}$. $\Gamma_* = 1$ indicates perfect prediction of the ground truth velocity, and $\Gamma_* = 0$ means the model has no predictive improvement over using the phase-averaged behavior. The data-driven models were notably more predictive than the template models, as illustrated in the right columns of Figure 4.1.

4.1.7 Discussion and Conclusions

We have shown that the broad class of "Shape-Underactuated Dissipative Systems (SUDS)" gives rise to dynamics that have an affine structure in the shape-velocity of their controlled DoF. As a consequence, it was possible for us to formulate an

---

[2] Equivalently, this can be considered a projection to the template system, which is a phase oscillator on the phase-averaged trajectory.

efficient regression model of these dynamics and to demonstrate that for several simple models, these regressions would in fact improve prediction accuracy by a substantial factor. The similarity to our previous work [BHR18; KBR19; BR20b] suggests that this would make it possible to rapidly learn gaits and libraries of maneuvers in such underactuated systems [BR19; BR20a]. It suggests that underactuation in SUDS does not pose nearly the same difficulties as in other underactuated systems — the strong dissipation improves the stability of the passive dynamics under repeated but perturbed control inputs.

One particularly promising direction is modeling and control of soft systems with e.g. soft pneumatic actuators or systems with long, passive, flexible tails. We have shown that our model identification regressions grow only linearly in complexity with the number of passive degrees of freedom. Thus, we can reasonably hope to process high dimensional representations of the continuous (and thus "infinite-dimensional") shape of soft objects. As long as the dimension of the representation provides a reliable state – in the sense of having good enough predictive ability – our work here provides good reason to believe the SUDS model identification will be tractable and produce predictive results.

From a biological perspective, we note that most animals are small (by human standards) and thus more dissipative because viscous friction scales with area or length, whereas inertia scales with volume. The simplicity of SUDS modeling suggests that the control problem that small, and even more so small and aquatic, animals solve is thus fundamentally easier than the control problem faced by large terrestrial creatures such as ourselves. We, therefore, offer the hypothesis that the neuromechanical control of animals is ancestrally geared for controlling SUDS and that the motor control ability of large-bodied extant species builds upon a more basal ability to learn to control SUDS.

A great part of the appeal of data-driven modeling to the robotics practitioner is the potential of our approach to systematically model the interactions of robots with un-modeled environments, even when these are potentially soft, compliant, and complex robots. Because the model regressions are efficient and easy to update, one can envision online identification leading to a broadly applicable form of adaptive control. This could allow robots to be highly adaptable to environmental changes and internal damage while retaining the ability to plan using the SUDS regression derived self-model.

Having provided a generalized framework for modeling shape-underactuated dissipative systems from data, we hope to inspire implementations in locomotion, manipulation, and even biomedical devices. For such applications, one needs to be sure of the dominance of damping and fairly high bandwidth control in a subspace of the shape of the robot. Having these, the practitioner has access to a system identifier that is sample efficient enough to work *in situ*, offering a broader space of practical applications for soft robots. These could include disaster scenarios with poorly characterized environments and biomedical procedures.

4.2 Modeling Extensions to the Perturbed Stokes Regime

4.2.1 Motivation

In this work, we provided tools for the modeling of animal and robot behaviors from data, specifically attending to the case where damping forces dominate inertial forces. A key consideration of this work was accounting for non-zero inertial forces, turning our attention to the perturbed Stokes regime [EJ16]. The geometric mechanics literature has historically paid very close attention to the dynamical regime where the inertia-damping ratio is zero [KM; KM95; HC11a; HC13; BHR18], resulting in the Stokesian limit of motion characterized by "living life at low Reynolds number" [Pur76]. This friction dominated dynamical regime can also result from nonholonomic constraints [Bre81; Kar81; Eld16].

Of specific interest in this work is the ability to model, analyze, and optimize animal and robot behaviors from data. The work of [BHR18] made progress in our ability to do this, combining tools from geometric mechanics and oscillator theory. A limiting feature of the methods was the required assumptions that the agent is restricted to act kinematically, restricted to the expressive capacity of the viscous connection [KM95]. One could argue that the method [BHR18] could be applied to systems with inertial effects, as long as the system acts approximately like a Stokesian system. This work showed that such an approach results in suboptimal modeling of system behaviors. In this work, we developed a data-driven modeling algorithm for systems inhabiting the perturbed Stokes regime, providing corrections to the model of [BHR18] that account for structural changes in the dynamics that happen when adding inertia to an otherwise fully viscous system. We showed that this model is more accurate than the viscous model for a range of inertia-damping ratios.

We arrived at a model for perturbed Stokesian behavior by applying normally hyperbolic invariant manifold theory (abbreviated NHIM theory) [Fen71; Fen74; HPS77; Fen79; Eld13] within a singular perturbation representation. Here the partitioning of fast and slow dynamics is governed by a scalar inertia-damping ratio. We showed that the slow manifold of this system acts like a viscous connection at a small inertia-damping ratio. At the lower limit of this ratio, we show that the system acts exactly like the Stokesian systems studied in [KM; KM95; HC11a; BHR18]. This representation is convenient since it provides a single constant that can scale us from purely viscous systems through the perturbed Stokes regime to what the geometric mechanics community calls the mixed kinematic dynamic regime. The representation has practical utility in that the singular perturbation representation allowed us to compute correction terms for the system dynamics when it is off of the slow manifold. This is a conventional use of geometric singular perturbation theory [Fen79; Jon95]. Researchers have used these methods in the past for robots with flexible joints [SKK87]. Please see subsections 2 and 3 of [KBR19] for details of the approach.

4.2.2 Designing Tests to Determine Performance Advantages of the Perturbed Stokes Model

After motivating a new set of regressors to plug into the modeling framework of [BHR18], we moved on to testing the utility of the new perturbed Stokes regressors. We designed a mechanical system for which an inertia-damping ratio can be prescribed as a scalar, allowing for the sampling of dynamics in the Stokesian, perturbed Stokesian, mixed kinematic dynamic, and inertial regimes. We showed that over a consistent region of the inertia-damping ratio, the perturbed Stokesian regressors offer predictive improvements over the Stokesian regressors. To be able to make this type of comparison, we introduced metrics for the evaluation of both model quality and comparative model performance. These metrics are computable from data, unlike the inertia-damping ratio, which is a prescribed variable. The methods here extended the sampling efficiency and morphology agnostic characteristics of [BHR18] to systems with momentum. The methods detailed in this work also increased the breadth of the approach of data-driven geometric modeling.

A key insight of this work is the derivation of a new collection of regressors based on the perturbed Stokesian dynamics, as characterized in [KBR19]. Here, we augment the modeling process of [BHR18] to use this new collection of regressors for estimating the motion model in the neighborhood of a behavior. This allows us to take an approach designed for systems acting at the limit of friction and extend to a class of systems with measurable momentum in the perturbed Stokes regime. In order to showcase the utility of this model, we simulate the dynamics of a collection of swimming platforms at various Reynolds numbers. We document the difference in the accuracy of the Stokes and perturbed Stokes data-driven models across these

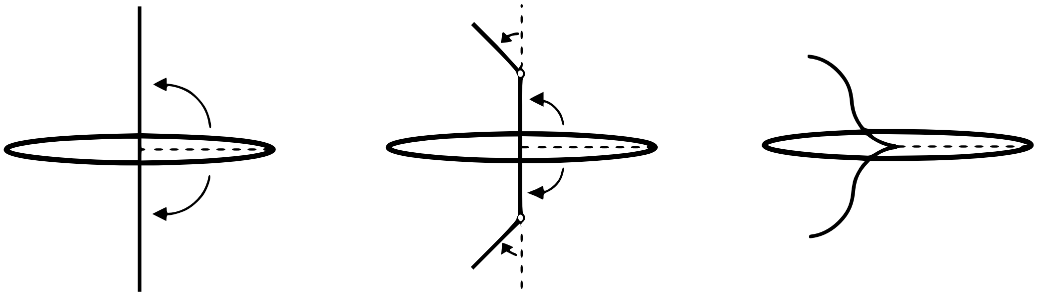

Figure 4.2: Schematic representation of our swimming model. A single body (ellipses with center of mass marked) of mass m and moment of inertia $m\bar{I}$ is attached to two identical paddles each comprising 1 (left), 2 (middle), or n segments (right). The length of the body is L, and the length of the paddles is d, with each segment of length d/n.

varying dynamical regimes. [3].

4.2.3 Designing a Swimming Platform that Scales through Dynamical Regimes

We designed a swimming model to test the accuracy of the two models. The system in Fig 4.2 has a central body (with uniformly distributed mass) and two chains of massless links, which we call paddles, that extend from the center of mass. Each paddle can be equally segmented into n pieces, which always sum to the same unit length $\frac{L}{2}$. At $n = 1$, the system reminds us of a boat with oars, whereas n approaches inf, the system reminds us of a bacterial cell with two flagella.

The system is constrained to move within a planar workspace, where the dynamics in the body frame are assumed to be invariant to the position and orientation in the plane. We represent this position and orientation on the plane as a group element $g \in \mathsf{SE}(2)$. This element is defined $g = [x, y, \theta]^T$ for the equations below.

---

[3] Are simulations are not intended to provide high fidelity characterization of fluid-fluid interactions at higher-Reynolds number. We did not focus on real-world physical fidelity at this range of Reynolds number since it is unlikely to showcase behavior we consider to be perturbed Stokesian.

The velocity in the body frame is defined as

$$\mathring{g} = R^{-1}(\theta)\dot{g} = \begin{bmatrix} \cos(\theta) & \sin(\theta) & 0 \\ -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \dot{g}. \tag{4.33}$$

The main body (length L) and the paddle links (length d) are treated as slender members. We use drag forces according to Cox theory [Cox70] to model the environmental forces on the system as

$$C_{\scriptscriptstyle D} = c \begin{bmatrix} C_x d & 0 & 0 \\ 0 & C_y d & 0 \\ 0 & 0 & \frac{1}{12} d^3 C_y \end{bmatrix}, \quad C_{\scriptscriptstyle L} = c \begin{bmatrix} C_x L & 0 & 0 \\ 0 & C_y L & 0 \\ 0 & 0 & \frac{1}{12} L^3 C_y \end{bmatrix}. \tag{4.34}$$

Factor $c > 0$ is written such that it can be used in our analyis as a scaling factor for the influence of drag in the system. According to Cox theory, at the limit of thinness in links (which we use here), the drag coefficient ratio can be as high as 2. [HC13]. Using these assumptions, we can write the drag forces and drag moments on the ith segment as

$$F_i = c\bar{F}_i - R(\alpha_i) C_{\scriptscriptstyle D} \begin{bmatrix} \cos(\alpha_i) & \sin(\alpha_i) & 0 & 0 \\ -\sin(\alpha_i) & \cos(\alpha_i) & \frac{d}{2} & \frac{d}{2} \\ 0 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} \mathring{g} \\ \dot{\alpha}_i \end{bmatrix}. \tag{4.35}$$

For the main body, we can write the drag force and drag moment as

$$F_{body} = c\bar{F}_{body} = -C_{\scriptscriptstyle L} \mathring{g}. \tag{4.36}$$

Provided that the body has a uniformly distributed mass m and moment of inertia $I = m\bar{I}$ about its midpoint (remember that the oars/segments are massless), the dynamics in the world frame are :

$$\ddot{g} = \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{\theta} \end{bmatrix} = \frac{1}{\epsilon} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{\bar{I}} \end{bmatrix} R(\theta) \left(\bar{F}_{d_1} + \bar{F}_{d_2} + \bar{F}_{body} \right), \tag{4.37}$$

where $\epsilon := \frac{m}{c}$ is the dimensionless ratio of inertia to damping. For dimension counting, remember that m and c are dimensionless, making the diagonal terms of equation 4.37 have units of $\frac{1}{time}$.

Thus, by adjusting the ratio of inertial to frictional forces in the system, we can directly modify ϵ. Now we can easily modify the regime the paddle operates in by dialing a single constant, ϵ. As ϵ approaches zero, the system acts in the fully viscous regime. Here the system is driftless; when the paddles stop, the system stops. As ϵ approaches inf, the system is governed by conservation of momentum. At some point, while increasing *epsilon*, the system will break the assumptions required to consider it a perturbed Stokesian system. The best way to see when this matters for our modeling assumptions is to test it. Next, we will compare the performance of data-driven models (Stokesian and perturbed Stokesian) at a variety of ϵ values that span the low to high Reynolds regimes.

4.2.4 Comparison of the Estimated Models

For the simulations run, we select drag coefficients $C_x = 1$ and $C_y = 2$, inertial coefficient $\bar{I} = 1$, and kinematic parameters $L = 1$ and $d = 0.5$. In the simulations, we will modify the value of ϵ, which governs the rate of attraction to the slow manifold as well as the ratio of inertial to viscous forces. Previous work [BHR18] documents the process for simulating and modeling systems from data, which is identical to the implementation used here. To summarize, we obtain a noisy trajectory in shape that is attracted to follow some nominal gait. This trajectory contains 30 cycles that loosely follow the nominal gait and are then integrated through equation 4.37.

The sampled outputs of the simulation are recorded and stored. The “ground truth” samples of the body velocity are referenced here as $\mathring{g}_G$. For each of these samples, we obtain a velocity that is forecasted with a data-driven perturbed Stokes model, which we call $\mathring{g}_p$. Likewise, we obtain velocities forecasted by the Stokes model, which we call $\mathring{g}_s$.

As a “zeroth-order” phase model of the dynamics, we constructed a Fourier series model of $\mathring{g}_G$ with respect to the estimated phase, which we denote by $\mathring{g}_a$. For any data point, the zeroth-order model prediction is $\mathring{g}_a(\varphi)$ for the phase φ of that data point.

We computed the RMS errors e_*^k for each component k of the body velocity and each model $* = p, s, a$ by $e_*^k := \langle |\mathring{g}_*^k - \mathring{g}_G^k|^2 \rangle^{1/2}$. Since the numerical value of these errors means little, we defined the metric $\Gamma_*^k := 1 - e_*^k / e_a^k$ for $* = p, s$ to indicate how much better the regression models were performing compared to the zeroth-order phase model $\mathring{g}_a$. A Γ_*^k of 0 indicates doing no better than the zeroth-order model, whereas a 1 indicates a perfect model. To further highlight the *difference* in prediction quality, we also plot $\Delta^k := \Gamma_p^k - \Gamma_s^k$.

4.2.4.1 Algorithm comparison using manually selected gaits

We chose to first test the modeling approaches on a collection of simple, manually selected behaviors. These include behaviors we term “twist in place” and “symmetric flapping” gaits, both of which initialize with paddles aligned at a quarter turn away from the body (as depicted in the two-segment model in Figure 4.2), and respectively involve anti-symmetric and symmetric sinusoidal movement of the paddles with amplitude 1. The “symmetric flapping gait” primarily moves in the direction of the x body axis, while the “twist in place gait” primarily changes the θ body coordinate. Finally, we considered a “circle” gait, which also initializes the paddles at a quarter

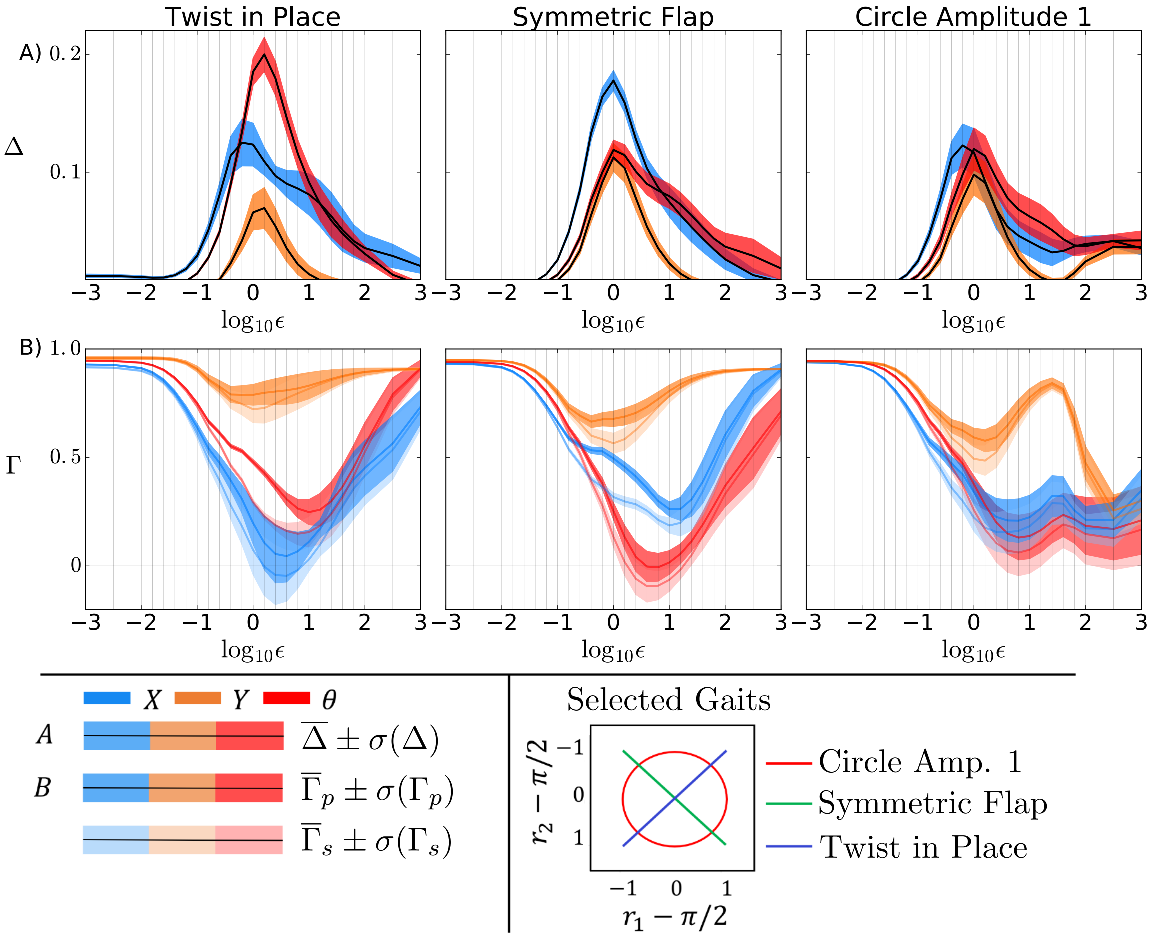

Figure 4.3: Comparison of model prediction quality when using the perturbed Stokes regressors versus the Stokes regressors on three gaits, in terms of the Γ and Δ quality metrics. We have plotted the components of Δ, representing the relative advantage of perturbed Stokes regressors (top row; (A)), and Γ, representing model prediction quality (bottom row; (B)), against 6 orders of magnitude variation in the inertial to viscosity ratio ϵ (logarithmic scale; sampled at 25 values (vertical gray lines). We present three gaits, whose shape space loci are in-phase paddle angle (which leads to anti-phase paddle motions; "Twist in Place"; left column; blue line in shape-space plot), anti-phase paddle angle (bilaterally symmetric paddle motions; "Symmetric Flap"; middle column; green line in shape-space plot), and quarter-cycle out of phase paddle angles ("Circle Amp. 1"; right column; red line in shape-space plot). All three gaits have paddle angles ranging between -1 and 1 radians. For each value of ϵ we performed 8 simulation trials each consisting of 30 (noisy) gait cycles, and plotted mean and standard deviation of Δ and Γ for each component of the $\mathfrak{se}(2)$ body motion (X blue; Y orange; θ red; saturated for Δ and Γ_p, pale for Γ_s). Consistently for all components and gaits, the perturbed Stokes regressors provide a better model for an order of magnitude or a wider range of ϵ around $\epsilon = 1$. For Twist in Place and Symmetric Flap gaits, both models are accurate for large and small ϵ (Γ close to 1); for the Circle Amplitude 1 gait, the prediction is only accurate for the Stokes regime (small ϵ).

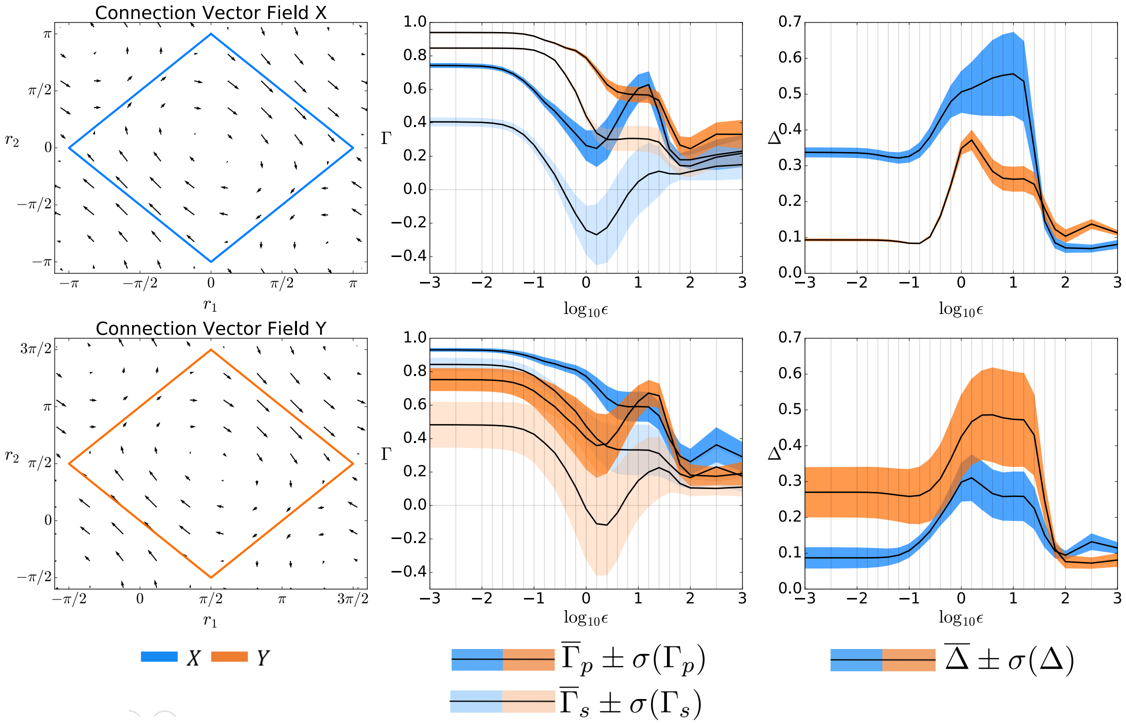

Figure 4.4: Comparison of model prediction quality when using the perturbed Stokes regressors versus the Stokes regressors on two extremal gaits, in terms of the Γ and Δ quality metrics. Plots consist of the same types as those in Fig 4.3. We only plot the X (blue) and Y (orange) components of Γ (middle column; saturated color Γ_p; pale colors Γ_s) and Δ (right column). We selected the gait to maximize either the X component of total body frame motion (top row) or the Y component (bottom row). The gaits are extremal in the Stoke regime ($\epsilon = 0$) and selected by taking the zero level set of the connection curvature (method from [HC11a; HC13]). Following their approach, we plot the connection of the coordinate being optimized as a vector field over the shape-space (black arrows; left column), with the shape-space gait locus plotted over it (diamond shapes in left column, colored by coordinate optimized). Results show that both models are most accurate for small ϵ (the Stokes regime; Γ closer to 1), with the perturbed Stokes regressors providing improvements across the entire range. Over the two order of magnitude range of $10^{-0.5} < \epsilon < 10^{1.5}$ this advantage is noticeably more pronounced (the perturbed Stokes regime; bump in Δ plots). Also note that X extremal gait shows much greater Δ^x; Y extremal gait shows much greater Δ^y.

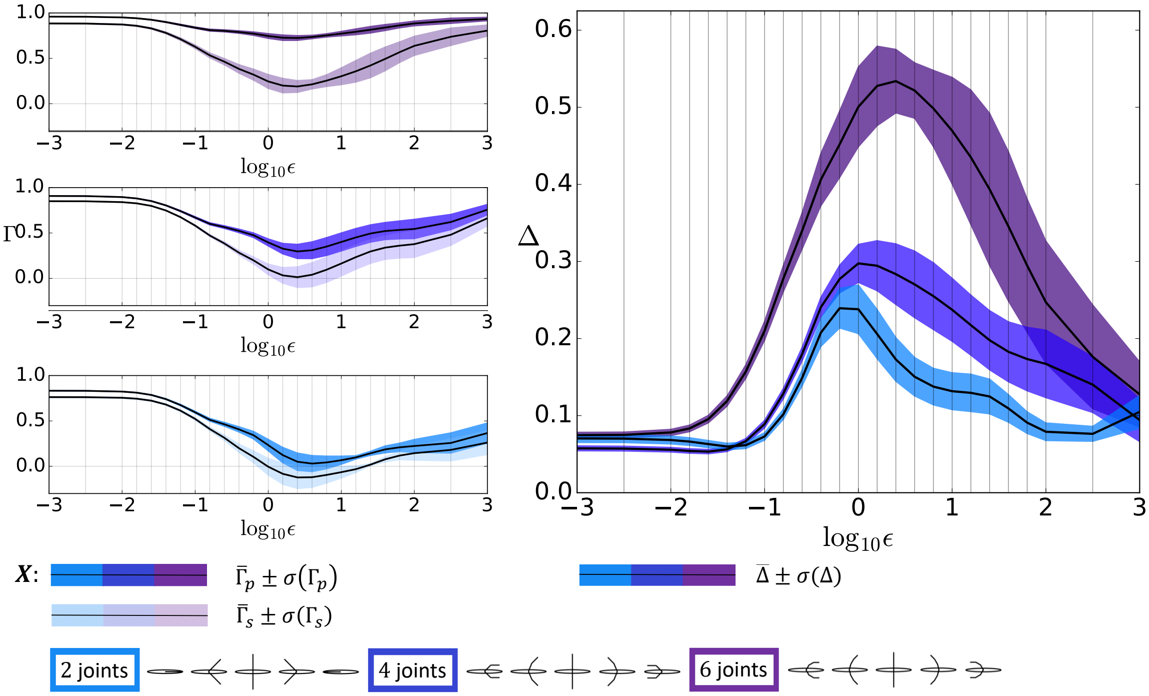

Figure 4.5: Comparison of model prediction quality when using the perturbed Stokes regressors versus the Stokes regressors on paddles with different dimensions shape space, shown in terms of the Γ and Δ quality metrics. Plots consist of the same types as those in Fig 4.3. We plotted Γ and Δ of three swimmers with different numbers of paddle segments: one segment per paddle (light blue), two segments (blue), and three segments (purple); see Fig 4.2 for schematic. We used a symmetric flapping gait (see Fig 4.3; small cartoons above). The paddles moved symmetrically with total angles of all joints summing up to a sinusoid of amplitude π. We plot the X components of Γ (left column; one plot per model; saturated colors Γ_p; pale colors Γ_s) and Δ (right column). Results show that over the two orders of magnitude range of $10^{-0.5} < \epsilon < 10^{1.5}$, the perturbed Stokes regressors consistently provide improvements. The relative improvement Δ increased markedly with shape space dimension, by as much as 0.5 in Δ.

turn away from the body and moves them sinusoidally with amplitude 1, but has a quarter cycle phase offset between them. This gait tends to move the system in a way that changes all three body coordinates throughout its execution.

We selected these three gaits because they are simple to describe and span a range of resultant body motions. For single link paddles, the body shape space is 2D, and these gaits are represented by loci that are diagonal lines with slopes 1, -1, and a circle (see Fig 4.3). We simulated the gaits and plotted mean and variance of Γ_s, Γ_p and Δ for each value of ϵ (Fig 4.3). The plot shows that for all three gaits tested and for all three body coordinates, over a range spanning an order of magnitude or more around $\epsilon = 1$, the perturbed Stokes models are better by $\Delta > 0.05$ or more.

4.2.4.2 Algorithm comparison using extremal gaits

Arbitrarily selected gaits such as those examined in the previous subsection are not expected to exhibit any special properties with respect to our modeling approach. In particular, with respect to a goal function $\phi(\cdot)$, they are expected to be regular points of $\phi(\cdot)$. However, ϕ-optimal gaits have $\nabla_p \phi = 0$ and thus have an additional structure that might interact with the modeling approach.

We chose goal functionals $\int \overset{\circ}{g}^x(t)\,\mathrm{d}t$ and $\int \overset{\circ}{g}^y(t)\,\mathrm{d}t$ (where superscripts denote components) corresponding to displacement in the x and y coordinates as measured in the body frame of the paddleboat. This is *not* the same as actual x or y displacement in the world, since boat orientation changes over time. Using the methods of [HC13], we determined the extremal gaits for these goal functionals in the Stokes regime with high accuracy. Plotted in the shape-space (and superimposed on the "connection vector fields" [HC11a; HC13] of the appropriate goal functional) they are diamond shaped (Fig 4.4). We also plotted Γ and Δ, revealing that again, perturbed Stokes regressors improve performance ($\Delta > 0.15$) over a range of two orders of magnitude in ϵ. Unlike the arbitrary gaits of the previous subsection, the extremal gaits have

$\Gamma > 0.1$ for all $\epsilon > 1$ for both model types. This suggests that even outside the perturbed Stokes regime, the addition of regressors improves upon the zeroth-order phase model. It is also notable that in the extremal x gait, Δ^x is significantly better than Δ^y, whereas in the extremal y gait, the converse is true.

4.2.4.3 Performance gains grow with shape space dimension

Thus far, we have only presented results for systems having 2D shape spaces. Because data-driven methods are often handicapped by their inability to scale with model dimensionality, we also chose to test our approach on systems of higher dimension by extending each paddle into a multi-segmented model. We selected a gait similar to that of the symmetric flapping gait, but with the additional feature that the bending angle of a paddle was uniformly distributed through the joints it contains. In particular, the relative angles between adjacent segments were equal and of amplitude π/N, where N is the number of joints.

We plotted Γ_p^x, Γ_s^x and Δ^x for paddles with 1, 2 and 3 segments (Fig 4.5). The Δ^x shows a marked improvement in the 4D and 6D models, suggesting that as shape-space complexity increased, the advantage of perturbed Stokes regressors became comparatively more significant.

4.2.5 Discussion and Conclusions

The results of [KBR19] showed that for variations of the platform morphology and gaits selected, there appears to be a sizable range of the inertia to damping ratio where the perturbed Stokes model is better at forecasting system behavior than the Stokes model. At $\epsilon = 0$, the system behavior collapses to that of the Stokesian regime. The metric Δ was introduced in [KBR19] to show the improvement in the predictive quality of the perturbed Stokes regressors with respect to the Stokes regressors. In particular, the improvement is consistently present in the region $\log_{10}\epsilon \in [0, 1]$, suggesting that this range of ϵ might be the range for which the predicted slow manifold is both present and sufficiently simple to be captured by the new regressors for this system.

The perturbed Stokes regressors seem to improve prediction performance more in the direction in which the gait was extremal. We hypothesize that this is because extremal gaits have already exhausted any first-order improvements available, i.e. gradients are zero. With the first-order terms close to zero, the presence of more high-order terms among the perturbed Stokes regressors may have a greater effect on the relative prediction error.

It is interesting to note the large magnitude of improvement in Δ as the shape space dimension increased in the paddle boat. Whether this is an artifact of the particular model and/or gait, or a more general feature, remains to be determined.

At the lower end ϵ magnitudes studied here, the systems are near the Stokesian limit, and therefore we expect relatively little improvement from adding regressors designed for the perturbed Stokes regime. This is consistent with our experimental results in all figures, which show for ϵ small both small values of Δ and large values of Γ for both sets of regressors.

There are at least two factors that prohibit the models from producing accurate models in the higher ranges of ϵ.

1. Terms of $\mathcal{O}(\epsilon^2)$ or higher have a greater impact on the dynamics at higher ϵ. One approach to having more accurate models at higher ϵ could be determining these higher-order terms and including them in the regression.

2. As ϵ increases, at some value the slow manifold will disappear. The inclusion of momentum states will become a necessity for capturing the dynamics of the system.

The novelty of the contribution in this work is the use of an approximation of behavior in the perturbed Stokes to inform an algorithm to generally model systems of this class from observational data. We show that this estimation is useful by demonstrating that it is more predictive than the Stokes-informed models at a sizable range of small inertia to damping ratios ϵ. We tested the perturbed Stokes models on the example system over a variety of gaits, obtaining improvements of $\Delta = 5$–65% compared to the Stokes models. Furthermore, the results of one of our experiments showed further improvements as the shape-space dimension of the locomoting system increased. This suggests that higher-dimensional systems might be modeled effectively using our approach.

4.2.6 Uses of Chapter Material in New Research

Future research will involve the implementation of this framework to optimize the behaviors of robots with sizable momentum. We can also begin to ask questions about the optimality of animal behaviors, even when these animals are not restricted to Stokesian behavior. This will allow us to look beyond the optimality of behaviors in systems such as *C. elegans*. We can now look at a broader range of systems that have sizable momentum or ϵ.

Some new research at the intersection of biology and robotics involves the classification of animal and robot behaviors as Stokesian, perturbed Stokesian, and mixed kinematic-dynamic. Metrics like Δ can aid in the distinction of these transitions since they can fairly compare which dynamical features are being expressed in the motion.

BIBLIOGRAPHY

[AB97] M. Ahmadi and M. Buehler. “Stable control of a simulated one-legged running robot with hip and leg compliance”. In: *IEEE Transactions on Robotics and Automation* 13.1 (1997), pp. 96–104. DOI: 10.1109/70.554350.

[Ana+18] Abhishek Ananthakrishnan et al. “Automated Gait Generation for Simulated Bodies Using Deep Reinforcement Learning”. In: *2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*. IEEE, Apr. 2018. DOI: 10.1109/icicct.2018.8473310.

[AR08] Joseph E Avron and Oren Raz. “A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin”. In: *New Journal of Physics* 10.6 (2008), p. 063016. DOI: 10.1088/1367-2630/10/6/063016.

[Ast+20] Henry C. Astley et al. “Surprising simplicities and syntheses in limbless self-propulsion in sand”. In: *The Journal of Experimental Biology* 223.5 (Feb. 2020), jeb103564. DOI: 10.1242/jeb.103564.

[Ban+16] Somil Bansal et al. “Learning quadrotor dynamics using neural network for flight control”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, Dec. 2016. DOI: 10.1109/cdc.2016.7798978.

[Baz+17] Salah Bazzi et al. “Motion analysis of two-link nonholonomic swimmers”. In: *Nonlinear Dynamics* 89.4 (June 2017), pp. 2739–2751. DOI: 10.1007/s11071-017-3622-y.

[BF93] R. Blickhan and R.J. Full. “Similarity in multilegged locomotion: Bouncing like a monopode”. In: *Journal of Comparative Physiology A* 173.5 (Nov. 1993). DOI: `10.1007/bf00197760`.

[BHR18] Brian Bittner, Ross L. Hatton, and Shai Revzen. “Geometrically optimal gaits: a data-driven approach”. In: *Nonlinear Dynamics* 94.3 (July 2018), pp. 1933–1948. DOI: `10.1007/s11071-018-4466-9`.

[BHR20] Brian Bittner, Ross L Hatton, and Shai Revzen. “Data-Driven Geometric System Identification for Shape-Underactuated Dissipative Systems”. In: *arXiv preprint arXiv:2012.11064* (2020).

[BLM] A.M. Bloch, N.E. Leonard, and J.E. Marsden. “Stabilization of mechanical systems using controlled Lagrangians”. In: *Proceedings of the 36th IEEE Conference on Decision and Control*. IEEE. DOI: `10.1109/cdc.1997.657135`.

[Blo+96] Anthony M Bloch et al. “Nonholonomic mechanical systems with symmetry”. In: *Archive for Rational Mechanics and Analysis* 136.1 (1996), pp. 21–99. DOI: `10.1007/BF02199365`.

[BR19] Brian Bittner and Shai Revzen. “Optimizing Gaits for Coverage on Lie Groups [ABSTRACT ONLY]”. In: *Dynamic Walking*. 2019. URL: `https://youtu.be/vowiZFfVBoM`.

[BR20a] Brian Bittner and Shai Revzen. “A Robot Made of Tree Branches Can Learn to Move in 15 Minutes [ABSTRACT ONLY]”. In: *Dynamic Walking*. 2020. URL: `https://youtu.be/iRBT3RkhdTU?t=3660`.

[BR20b] Brian Bittner and Shai Revzen. “Geometric Gait Optimization with a five-link wheeled snake”. In: *Bulletin of the American Physical Society* (2020).

[Bre81] V.N. Brendelev. “On the realization of constraints in nonholonomic mechanics”. In: *Journal of Applied Mathematics and Mechanics* 45.3 (Jan. 1981), pp. 351–355. DOI: `10.1016/0021-8928(81)90065-4`.

[BRS15] Samuel A. Burden, Shai Revzen, and S. Shankar Sastry. “Model Reduction Near Periodic Orbits of Hybrid Dynamical Systems”. In: *IEEE Transactions on Automatic Control* 60.10 (Oct. 2015), pp. 2626–2639. DOI: `10.1109/tac.2015.2411971`.

[Bru+20] Daniel Bruder et al. “Data-Driven Control of Soft Robots Using Koopman Operator Theory”. In: *IEEE Transactions on Robotics* (2020), pp. 1–14. DOI: `10.1109/tro.2020.3038693`.

[BZL06] J. Bongard, V. Zykov, and H. Lipson. “Resilient Machines Through Continuous Self-Modeling”. In: *Science* 314.5802 (Nov. 2006), pp. 1118–1121. DOI: `10.1126/science.1133687`.

[Cal+14] Roberto Calandra et al. “An experimental comparison of Bayesian optimization for bipedal locomotion”. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2014. DOI: `10.1109/icra.2014.6907117`.

[CHG20] G. T. Clifton, D. Holway, and N. Gravish. “Uneven substrates constrain walking speed in ants through modulation of stride frequency more than stride length”. In: *Royal Society Open Science* 7.3 (Mar. 2020), p. 192068. DOI: `10.1098/rsos.192068`.

[CMR01] Hernán Cendra, Jerrold E. Marsden, and Tudor S. Ratiu. “Geometric Mechanics, Lagrangian Reduction, and Nonholonomic Systems”. In: *Mathematics Unlimited — 2001 and Beyond*. Springer Berlin Heidelberg, 2001, pp. 221–273. DOI: `10.1007/978-3-642-56478-9_10`.

[Cox70] R. G. Cox. “The motion of long slender bodies in a viscous fluid Part 1. General theory”. In: *Journal of Fluid Mechanics* 44.04 (Dec. 1970), p. 791. DOI: `10.1017/s002211207000215x`.

[CP20] Yu-Ming Chen and Michael Posa. “Optimal Reduced-order Modeling of Bipedal Locomotion”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2020. DOI: `10.1109/icra40945.2020.9197004`.

[CR20] George Council and Shai Revzen. “Fast Recovery of Robot Behaviors”. In: *arXiv preprint arXiv:2005.00506* (2020).

[Da+16] Xingye Da et al. “From 2D Design of Underactuated Bipedal Gaits to 3D Implementation: Walking With Speed Tracking”. In: *IEEE Access* 4 (2016), pp. 3469–3478. DOI: `10.1109/access.2016.2582731`.

[Dai+] Jin Dai et al. “Geometric Swimming on a Granular Surface”. In: *Robotics: Science and Systems XII*. Robotics: Science and Systems Foundation. DOI: `10.15607/rss.2016.xii.012`.

[Dea+20] Tony Dear et al. “Locomotion of a multi-link non-holonomic snake robot with passive joints”. In: *The International Journal of Robotics Research* 39.5 (Jan. 2020), pp. 598–616. DOI: `10.1177/0278364919898503`.

[DR11] Marc Deisenroth and Carl E Rasmussen. “PILCO: A model-based and data-efficient approach to policy search”. In: *Proceedings of the 28th International Conference on machine learning (ICML-11)*. 2011, pp. 465–472. URL: `http://aiweb.cs.washington.edu/research/projects/aiweb/media/papers/tmpZj4RyS.pdf`.

[EJ16] Jaap Eldering and Henry O. Jacobs. “The Role of Symmetry and Dissipation in Biolocomotion”. In: *SIAM Journal on Applied Dynamical Systems* 15.1 (Jan. 2016), pp. 24–59. DOI: `10.1137/140970914`.

[Eld13] Jaap Eldering. *Normally Hyperbolic Invariant Manifolds*. Atlantis Press, 2013. DOI: 10.2991/978-94-6239-003-4.

[Eld16] Jaap Eldering. "Realizing nonholonomic dynamics as limit of friction forces". In: *Regular and Chaotic Dynamics* 21.4 (July 2016), pp. 390–409. DOI: 10.1134/s156035471604002x.

[FDF05] E. Frazz, M.A. Dahleh, and E. Feron. "Maneuver-based motion planning for nonlinear systems with symmetries". In: *IEEE Transactions on Robotics* 21.6 (Dec. 2005), pp. 1077–1091. DOI: 10.1109/tro.2005.852260.

[Fen71] Neil Fenichel. "Persistence and Smoothness of Invariant Manifolds for Flows". In: *Indiana University Mathematics Journal* 21.3 (1971), pp. 193–226. DOI: 10.1512/iumj.1972.21.21017.

[Fen74] N Fenichel. "Asymptotic stability with rate conditions". In: *Indiana University Mathematics Journal* 23.12 (1974), pp. 1109–1137. URL: https://www.jstor.org/stable/24890716?seq=1.

[Fen79] Neil Fenichel. "Geometric singular perturbation theory for ordinary differential equations". In: *Journal of Differential Equations* 31.1 (Jan. 1979), pp. 53–98. DOI: 10.1016/0022-0396(79)90152-9.

[Fit+17] Ian Fitzner et al. "Rapidly Prototyping Robots: Using Plates and Reinforced Flexures". In: *IEEE Robotics & Automation Magazine* 24.1 (Mar. 2017), pp. 41–47. DOI: 10.1109/mra.2016.2639058.

[FK99] Robert J Full and Daniel E Koditschek. "Templates and anchors: neuromechanical hypotheses of legged locomotion on land". In: *Journal of experimental biology* 202.23 (1999), pp. 3325–3332. URL: https://jeb.biologists.org/content/202/23/3325.

[FMJ02] Ajo Fod, Maja J Matarić, and Odest Chadwicke Jenkins. "Automated derivation of primitives for movement classification". In: *Autonomous robots* 12.1 (2002), pp. 39–54. DOI: 10.1023/A:1013254724861.

[Gan+15] Zhenyu Gan et al. "Passive Dynamics Explain Quadrupedal Walking, Trotting, and Tölting". In: *Journal of Computational and Nonlinear Dynamics* 11.2 (Aug. 2015). DOI: 10.1115/1.4030622.

[Gan+18] Zhenyu Gan et al. "All common bipedal gaits emerge from a single passive model". In: *Journal of The Royal Society Interface* 15.146 (Sept. 2018), p. 20180455. DOI: 10.1098/rsif.2018.0455.

[GG17] Brent Griffin and Jessy Grizzle. "Nonholonomic virtual constraints and gait optimization for robust walking control". In: *The International Journal of Robotics Research* 36.8 (May 2017), pp. 895–922. DOI: 10.1177/0278364917708249.

[GGC] Chaohui Gong, Daniel I. Goldman, and Howie Choset. "Simplifying Gait Design via Shape Basis Optimization". In: *Robotics: Science and Systems XII*. Robotics: Science and Systems Foundation. DOI: 10.15607/rss.2016.xii.006.

[GH83] John Guckenheimer and Philip Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer New York, 1983. DOI: 10.1007/978-1-4612-1140-2. URL: https://doi.org/10.1007%2F978-1-4612-1140-2.

[Gol+99] Martin Golubitsky et al. "Symmetry in locomotor central pattern generators and animal gaits". In: *Nature* 401.6754 (Oct. 1999), pp. 693–695. DOI: 10.1038/44416.

[GP07] Antoine Girard and George J. Pappas. “Approximate bisimulation relations for constrained linear systems”. In: *Automatica* 43.8 (Aug. 2007), pp. 1307–1317. DOI: `10.1016/j.automatica.2007.01.019`.

[Gre+14] Robert D. Gregg et al. “Evidence for a Time-Invariant Phase Variable in Human Ankle Control”. In: *PLoS ONE* 9.2 (Feb. 2014). Ed. by Manabu Sakakibara, e89163. DOI: `10.1371/journal.pone.0089163`.

[Gu+17] Shixiang Gu et al. “Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2017. DOI: `10.1109/icra.2017.7989385`.

[Ha+20] Sehoon Ha et al. “Learning to Walk in the Real World with Minimal Human Effort”. In: *arXiv preprint arXiv:2002.08550* (2020).

[Haa+18] Tuomas Haarnoja et al. “Composable Deep Reinforcement Learning for Robotic Manipulation”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2018. DOI: `10.1109/icra.2018.8460756`.

[Hat+13] Ross L. Hatton et al. “Geometric Visualization of Self-Propulsion in a Complex Medium”. In: *Physical Review Letters* 110.7 (Feb. 2013). DOI: `10.1103/physrevlett.110.078101`.

[Hau+08] Kris Hauser et al. “Using Motion Primitives in Probabilistic Sample-Based Planning for Humanoid Robots”. In: *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, 2008, pp. 507–522. DOI: `10.1007/978-3-540-68405-3_32`.

[HC10] Ross L Hatton and Howie Choset. “Optimizing coordinate choice for locomoting systems”. In: *2010 IEEE International Conference on Robotics and Automation*. IEEE, May 2010. DOI: `10.1109/robot.2010.5509876`.

[HC11a] Ross L Hatton and Howie Choset. “Geometric motion planning: The local connection, Stokes’ theorem, and the importance of coordinate choice”. In: *The International Journal of Robotics Research* 30.8 (June 2011), pp. 988–1014. DOI: 10.1177/0278364910394392.

[HC11b] Ross L Hatton and Howie Choset. “Geometric motion planning: The local connection, Stokes’ theorem, and the importance of coordinate choice”. In: *The International Journal of Robotics Research* 30.8 (June 2011), pp. 988–1014. DOI: 10.1177/0278364910394392.

[HC13] Ross L. Hatton and Howie Choset. “Geometric Swimming at Low and High Reynolds Numbers”. In: *IEEE Transactions on Robotics* 29.3 (June 2013), pp. 615–624. DOI: 10.1109/tro.2013.2251211.

[HC15] R.L. Hatton and H. Choset. “Nonconservativity and noncommutativity in locomotion”. In: *The European Physical Journal Special Topics* 224.17-18 (Dec. 2015), pp. 3141–3174. DOI: 10.1140/epjst/e2015-50085-y.

[HC85] Ronald M Harris-Warrick and Avis H Cohen. “Serotonin modulates the central pattern generator for locomotion in the isolated lamprey spinal cord”. In: *Journal of Experimental Biology* 116.1 (1985), pp. 27–46. URL: https://jeb.biologists.org/content/116/1/27.

[HDC17] Ross L. Hatton, Tony Dear, and Howie Choset. “Kinematic Cartography and the Efficiency of Viscous Swimming”. In: *IEEE Transactions on Robotics* 33.3 (June 2017), pp. 523–535. DOI: 10.1109/tro.2017.2653810. URL: https://doi.org/10.1109%2Ftro.2017.2653810.

[HDG17] Ross Hartley, Xingye Da, and Jessy W. Grizzle. “Stabilization of 3D underactuated biped robots: Using posture adjustment and gait libraries to reject velocity disturbances”. In: *2017 IEEE Conference on Control*

Technology and Applications (CCTA). IEEE, Aug. 2017. DOI: 10.1109/ccta.2017.8062632.

[Hee+17] Nicolas Heess et al. “Emergence of locomotion behaviours in rich environments”. In: *arXiv preprint arXiv:1707.02286* (2017).

[HPS77] Morris W. Hirsch, Charles C. Pugh, and Michael Shub. *Invariant Manifolds*. Springer Berlin Heidelberg, 1977. DOI: 10.1007/bfb0092042.

[Jon95] Christopher K. R. T. Jones. “Geometric singular perturbation theory”. In: *Dynamical Systems*. Springer Berlin Heidelberg, 1995, pp. 44–118. DOI: 10.1007/bfb0095239.

[Kaj+] S. Kajita et al. “Biped walking pattern generation by using preview control of zero-moment point”. In: *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*. IEEE. DOI: 10.1109/robot.2003.1241826.

[Kar81] A.V. Karapetian. “On realizing nonholonomic constraints by viscous friction forces and celtic stones stability”. In: *Journal of Applied Mathematics and Mechanics* 45.1 (Jan. 1981), pp. 30–36. DOI: 10.1016/0021-8928(81)90006-x.

[KBR19] Matthew D. Kvalheim, Brian Bittner, and Shai Revzen. “Gait modeling and optimization for the perturbed Stokes regime”. In: *Nonlinear Dynamics* 97.4 (Aug. 2019), pp. 2249–2270. DOI: 10.1007/s11071-019-05121-3.

[KCH99] G.K. Klute, J.M. Czerniecki, and B. Hannaford. “McKibben artificial muscles: pneumatic actuators with biomechanical intelligence”. In: *1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (Cat. No.99TH8399)*. IEEE, 1999. DOI: 10.1109/aim.1999.803170.

[KM] S.D. Kelly and R.M. Murray. “The geometry and control of dissipative systems”. In: *Proceedings of 35th IEEE Conference on Decision and Control*. IEEE. DOI: 10.1109/cdc.1996.574612.

[KM09] R.A. Knepper and M.T. Mason. “Path diversity is only part of the problem”. In: *2009 IEEE International Conference on Robotics and Automation*. IEEE, May 2009. DOI: 10.1109/robot.2009.5152696.

[KM95] Scott D. Kelly and Richard M. Murray. “Geometric phases and robotic locomotion”. In: *Journal of Robotic Systems* 12.6 (June 1995), pp. 417–431. DOI: 10.1002/rob.4620120607.

[KRC15] Simon Kalouche, David Rollinson, and Howie Choset. “Modularity for maximum mobility and manipulation: Control of a reconfigurable legged robot with series-elastic actuators”. In: *2015 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, Oct. 2015. DOI: 10.1109/ssrr.2015.7442943.

[Kva09] Matt Kvalheim. “Aspects of Invariant Manifold Theory and Applications”. PhD thesis. University of Michigan, Ann Arbor, 2009.

[Lau11] Eric Lauga. “Life around the scallop theorem”. In: *Soft Matter* 7.7 (2011), pp. 3060–3065. DOI: 10.1039/c0sm00953a.

[Mae+18] Azumi Maekawa et al. “Improvised Robotic Design with Found Objects [ABSTRACT ONLY]”. In: *Conference on Neural Information Processing Systems, Creativity Workshop* (2018). URL: https://azumi-maekawa.com/pdf/improvised_robotic_design_NeurIPS2018WS.pdf.

[Mah36] P.C. Mahalanobis. “On the generalized distance in statistics”. In: *Proc Natl Inst Sci India* 2.1 (1936), pp. 49–55. URL: http://library.isical.ac.in:8080/jspui/bitstream/123456789/6765/1/Vol02_1936_1_Art05-pcm.pdf.

[Mau+15] Horst-Moritz Maus et al. "Constructing predictive models of human running". In: *Journal of The Royal Society Interface* 12.103 (Feb. 2015), p. 20140899. DOI: 10.1098/rsif.2014.0899.

[McI+16] B. McInroe et al. "Tail use improves performance on soft substrates in models of early vertebrate land locomotors". In: *Science* 353.6295 (July 2016), pp. 154–158. DOI: 10.1126/science.aaf0984.

[MJD18] Jan Matas, Stephen James, and Andrew J Davison. "Sim-to-real reinforcement learning for deformable object manipulation". In: *arXiv preprint arXiv:1806.07851* (2018).

[MO98] Jerrold E Marsden and Jim Ostrowski. "Symmetries in motion: Geometric foundations of motion control". In: (1998). URL: https://core.ac.uk/download/pdf/4885891.pdf.

[Nag+18] Anusha Nagabandi et al. "Neural Network Dynamics for Model-Based Deep Reinforcement Learning with Model-Free Fine-Tuning". In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2018. DOI: 10.1109/icra.2018.8463189.

[OB98] Jim Ostrowski and Joel Burdick. "The Geometric Mechanics of Undulatory Robotic Locomotion". In: *The International Journal of Robotics Research* 17.7 (July 1998), pp. 683–701. DOI: 10.1177/027836499801700701.

[Ost+] J. Ostrowski et al. "Nonholonomic mechanics and locomotion: the snakeboard example". In: *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*. IEEE Comput. Soc. Press. DOI: 10.1109/robot.1994.351153.

[Ost96] James Patrick Ostrowski. "The mechanics and control of undulatory robotic locomotion". PhD thesis. California Institute of Technology, 1996.

[Ost99] J.P. Ostrowski. “Computing reduced equations for robotic systems with constraints and symmetries”. In: *IEEE Transactions on Robotics and Automation* 15.1 (1999), pp. 111–123. DOI: 10.1109/70.744607.

[Par+20] Clara Park et al. “An organosynthetic dynamic heart model with enhanced biomimicry guided by cardiac diffusion tensor imaging”. In: *Science Robotics* 5.38 (Jan. 2020), eaay9106. DOI: 10.1126/scirobotics.aay9106.

[Pel00] R. Pelrine. “High-Speed Electrically Actuated Elastomers with Strain Greater Than 100%”. In: *Science* 287.5454 (Feb. 2000), pp. 836–839. DOI: 10.1126/science.287.5454.836.

[PK11] Mihail Pivtoraiko and Alonzo Kelly. “Kinodynamic motion planning with state lattice motion primitives”. In: *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Sept. 2011. DOI: 10.1109/iros.2011.6094900.

[PKK09] Mihail Pivtoraiko, Ross A. Knepper, and Alonzo Kelly. “Differentially constrained mobile robot motion planning in state lattices”. In: *Journal of Field Robotics* 26.3 (Mar. 2009), pp. 308–333. DOI: 10.1002/rob.20285.

[Pur76] E. M. Purcell. “Life at low Reynolds number”. In: *AIP Conference Proceedings*. AIP, 1976. DOI: 10.1063/1.30370.

[PW] G.A. Pratt and M.M. Williamson. “Series elastic actuators”. In: *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*. IEEE Comput. Soc. Press. DOI: 10.1109/iros.1995.525827.

[Rad+15] Nicolaus A Radford et al. “Valkyrie: Nasa’s first bipedal humanoid robot”. In: *Journal of Field Robotics* 32.3 (2015), pp. 397–419.

[Rev09] Shai Revzen. "Neuromechanical control architectures of arthropod locomotion". PhD thesis. UC Berkeley, 2009.

[RG] Sam Roweis and Zoubin Ghahramani. "A Unifying Review of Linear Gaussian Models". In: *Neural Computation* 11.2 (), pp. 305–345. DOI: 10.1162/089976699300016674.

[RG08] Shai Revzen and John M. Guckenheimer. "Estimating the phase of synchronized oscillators". In: *Physical Review E* 78.5 (Nov. 2008). DOI: 10.1103/physreve.78.051907.

[RG11] Shai Revzen and John M. Guckenheimer. "Finding the dimension of slow dynamics in a rhythmic system". In: *Journal of The Royal Society Interface* 9.70 (Sept. 2011), pp. 957–971. DOI: 10.1098/rsif.2011.0431. URL: https://doi.org/10.1098%2Frsif.2011.0431.

[RH16] Suresh Ramasamy and Ross L. Hatton. "Soap-bubble optimization of gaits". In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, Dec. 2016. DOI: 10.1109/cdc.2016.7798407.

[RH17] Suresh Ramasamy and Ross L. Hatton. "Geometric gait optimization beyond two dimensions". In: *2017 American Control Conference (ACC)*. IEEE, May 2017. DOI: 10.23919/acc.2017.7963025.

[RH19] Suresh Ramasamy and Ross L. Hatton. "The Geometry of Optimal Gaits for Drag-Dominated Kinematic Systems". In: *IEEE Transactions on Robotics* 35.4 (Aug. 2019), pp. 1014–1033. DOI: 10.1109/tro.2019.2915424.

[RH20] Suresh Ramasamy and Ross L Hatton. "Optimal Gaits for Drag-dominated Swimmers with Passive Elastic Joints". In: *arXiv preprint arXiv:2010.01121* (2020).

[Ric+17] Leonardo Ricotti et al. “Biohybrid actuators for robotics: A review of devices actuated by living cells”. In: *Science Robotics* 2.12 (Nov. 2017), eaaq0495. DOI: `10.1126/scirobotics.aaq0495`.

[Rie+19] Jennifer M Rieser et al. “Geometric phase and dimensionality reduction in locomoting living systems”. In: *arXiv preprint arXiv:1906.11374* (2019).

[RK15a] Shai Revzen and Matthew Kvalheim. “Data driven models of legged locomotion”. In: *Micro- and Nanotechnology Sensors, Systems, and Applications VII*. Ed. by Thomas George, Achyut K. Dutta, and M. Saif Islam. SPIE, May 2015. DOI: `10.1117/12.2178007`.

[RK15b] Shai Revzen and Matthew Kvalheim. “Data driven models of legged locomotion”. In: *Micro- and Nanotechnology Sensors, Systems, and Applications VII*. Ed. by Thomas George, Achyut K. Dutta, and M. Saif Islam. SPIE, May 2015. DOI: `10.1117/12.2178007`.

[RKF09] Shai Revzen, Daniel E. Koditschek, and Robert J. Full. “Towards Testable Neuromechanical Control Architectures for Running”. In: *Advances in Experimental Medicine and Biology*. Springer US, 2009, pp. 25–55. DOI: `10.1007/978-0-387-77064-2_3`.

[RMH14] Elliott J. Rouse, Luke M. Mooney, and Hugh M. Herr. “Clutchable series-elastic actuator: Implications for prosthetic knee design”. In: *The International Journal of Robotics Research* 33.13 (Oct. 2014), pp. 1611–1625. DOI: `10.1177/0278364914545673`.

[Rol+14] David Rollinson et al. “Design and architecture of a series elastic snake robot”. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Sept. 2014. DOI: `10.1109/iros.2014.6943219`.

[RT15] Daniela Rus and Michael T. Tolley. “Design, fabrication and control of soft robots”. In: *Nature* 521.7553 (May 2015), pp. 467–475. DOI: `10.1038/nature14543`.

[RTS65] H. E. Rauch, F. Tung, and C. T. Striebel. “Maximum likelihood estimates of linear dynamic systems”. In: *AIAA Journal* 3.8 (Aug. 1965), pp. 1445–1450. DOI: `10.2514/3.3166`.

[Sas13] Shankar Sastry. *Nonlinear systems: analysis, stability, and control.* Vol. 10. Springer Science & Business Media, 2013. DOI: `10.1007/978-1-4757-3108-8`.

[SBC15] Samuel Shian, Katia Bertoldi, and David R. Clarke. “Dielectric Elastomer Based “Grippers” for Soft Robotics”. In: *Advanced Materials* 27.43 (Sept. 2015), pp. 6814–6819. DOI: `10.1002/adma.201503078`.

[SBK01] Uluc Saranli, Martin Buehler, and Daniel E. Koditschek. “RHex: A Simple and Highly Mobile Hexapod Robot”. In: *The International Journal of Robotics Research* 20.7 (July 2001), pp. 616–631. DOI: `10.1177/02783640122067570`.

[Sch+05] Stefan Schaal et al. “Learning Movement Primitives”. In: *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, 2005, pp. 561–572. DOI: `10.1007/11008941_60`.

[Sch+19] Perrin E. Schiebel et al. “Mechanical diffraction reveals the role of passive dynamics in a slithering snake”. In: *Proceedings of the National Academy of Sciences* 116.11 (Feb. 2019), pp. 4798–4803. DOI: `10.1073/pnas.1808675116`.

[Sei+17] Justin Seipel et al. “Conceptual Models of Legged Locomotion”. In: *Bioinspired Legged Locomotion*. Elsevier, 2017, pp. 55–131. DOI: `10.1016/b978-0-12-803766-9.00004-x`.

[Seo+13] Sangok Seok et al. “Meshworm: A Peristaltic Soft Robot With Antagonistic Nickel Titanium Coil Actuators”. In: *IEEE/ASME Transactions on Mechatronics* 18.5 (Oct. 2013), pp. 1485–1497. DOI: 10.1109/tmech.2012.2204070.

[SH00a] John Schmitt and Philip Holmes. “Mechanical models for insect locomotion: dynamics and stability in the horizontal plane - II. Application”. In: *Biological Cybernetics* 83.6 (Nov. 2000), pp. 517–527. DOI: 10.1007/s004220000180.

[SH00b] John Schmitt and Philip Holmes. “Mechanical models for insect locomotion: dynamics and stability in the horizontal plane I. Theory”. In: *Biological Cybernetics* 83.6 (Nov. 2000), pp. 501–515. DOI: 10.1007/s004220000181.

[SKK87] M. Spong, K. Khorasani, and P. Kokotovic. “An integral manifold approach to the feedback control of flexible joint robots”. In: *IEEE Journal on Robotics and Automation* 3.4 (Aug. 1987), pp. 291–300. DOI: 10.1109/jra.1987.1087102.

[SM92] S.S. Sastry and R. Montgomery. “The Structure of Optimal Controls for a Steering Problem”. In: vol. 25. 13. Elsevier BV, June 1992, pp. 135–140. DOI: 10.1016/s1474-6670(17)52270-3.

[Sün+18] Niko Sünderhauf et al. “The limits and potentials of deep learning for robotics”. In: *The International Journal of Robotics Research* 37.4-5 (Apr. 2018), pp. 405–420. DOI: 10.1177/0278364918770733.

[SW89a] Alfred Shapere and Frank Wilczek. “Efficiencies of self-propulsion at low Reynolds number”. In: *Journal of Fluid Mechanics* 198.-1 (Jan. 1989), p. 587. DOI: 10.1017/s0022112089000261.

[SW89b] Alfred Shapere and Frank Wilczek. “Geometry of self-propulsion at low Reynolds number”. In: *Journal of Fluid Mechanics* 198.-1 (Jan. 1989), p. 557. DOI: 10.1017/s002211208900025x.

[Tan+18] Jie Tan et al. “Sim-to-real: Learning agile locomotion for quadruped robots”. In: *arXiv preprint arXiv:1804.10332* (2018).

[TH07] Daniel Tam and A. E. Hosoi. “Optimal Stroke Patterns for Purcell’s Three-Link Swimmer”. In: *Physical Review Letters* 98.6 (Feb. 2007). DOI: 10.1103/physrevlett.98.068105.

[Tod11] Emanuel Todorov. “A convex, smooth and invertible contact model for trajectory optimization”. In: *2011 IEEE International Conference on Robotics and Automation*. IEEE, May 2011. DOI: 10.1109/icra.2011.5979814.

[Tol+14] Michael T Tolley et al. “An untethered jumping soft robot”. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2014, pp. 561–566. DOI: 10.1109/IROS.2014.6942615.

[Wal54] W. Grey Walter. “THE LIVING BRAIN”. In: *The Journal of Nervous and Mental Disease* 120.1 (July 1954), p. 144. DOI: 10.1097/00005053-195407000-00069.

[Wei+02] Joel D Weingarten et al. “Gait generation and optimization for legged robots”. In: *The IEEE International Conference on Robotics and Automation*. 2002. URL: https://kodlab.seas.upenn.edu/uploads/Kod/Weingarten03.pdf.

[Wes+18] Eric R Westervelt et al. *Feedback control of dynamic bipedal robot locomotion*. CRC press, 2018. URL: https://ece.umich.edu/faculty/grizzle/images/web-book/Fdbk_Ctrl_of_Dyn_Bipedal_Robot_Locomotion_excerpts.pdf.

[Wil+17] Simon Wilshin et al. “Longitudinal quasi-static stability predicts changes in dog gait on rough terrain”. In: *The Journal of Experimental Biology* 220.10 (Mar. 2017), pp. 1864–1874. DOI: 10.1242/jeb.149112.

[WJ10] Robert J. Webster and Bryan A. Jones. “Design and Kinematic Modeling of Constant Curvature Continuum Robots: A Review”. In: *The International Journal of Robotics Research* 29.13 (June 2010), pp. 1661–1683. DOI: 10.1177/0278364910368147. URL: https://doi.org/10.1177%2F0278364910368147.

[WL13] Loïc Was and Eric Lauga. “Optimal propulsive flapping in Stokes flows”. In: *Bioinspiration & Biomimetics* 9.1 (Dec. 2013), p. 016001. DOI: 10.1088/1748-3182/9/1/016001.

[WO16] Oren Wiezel and Yizhar Or. “Using optimal control to obtain maximum displacement gait for purcell’s three-link swimmer”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE. 2016, pp. 4463–4468. DOI: 10.1109/cdc.2016.7798947.

[WZR19] Ziyou Wu, Dan Zhao, and Shai Revzen. “Coulomb Friction Crawling Model Yields Linear Force–Velocity Profile”. In: *Journal of Applied Mechanics* 86.5 (Mar. 2019). DOI: 10.1115/1.4042696.