AN ABSTRACT OF THE THESIS OF

Suresh Ramasamy for the degree of Doctor of Philosophy in Robotics and Mechanical Engineering presented on August 26, 2020.

Title: Geometry of Locomotion

Abstract approved: _____

Ross L. Hatton

There are four major contributions described in this thesis:

1. A set of geometric principles for understanding and optimizing the gaits of drag-dominated kinematic locomoting systems. For systems with two shape variables, the dynamics of gait optimization are analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble. The internal pressure on the gait curve is provided by the flux of the curvature of the system constraints passing through the surface bounded by the gait, and surface tension is provided by the cost associated with executing the gait, which when executed at optimal (constant-power) pacing is proportional to its pathlength measured under a Riemannian metric. We extend these principles to work on systems with three, and then more than three, shape variables. We demonstrate these principles on a variety of system geometries (including Purcell's swimmer) and for optimization

criteria that include maximizing displacement and efficiency of motion for both translation and turning motions. We also demonstrate how these principles can be used to simultaneously optimize a system's gait kinematics and physical design.

2. Optimal swimming strategies for drag-dominated swimmers with a passive elastic joint. We use frequency domain analysis to relate the motion of the passive joint to the motion of the actuated joint. We couple this analysis with elements of the geometric framework introduced in our first contribution, to identify speed-maximizing and efficiency-maximizing gaits for drag-dominated swimmers with a passive elastic joint.
3. The first two contributions rely heavily on the total Lie bracket formula to identify useful gaits. The total Lie bracket formula is obtained using a truncated Baker-Campbell-Hausdorff (BCH) expansion, and its accuracy depends on the choice of coordinates used to describe the system [34]. We relate the magnitude of the higher order BCH terms to the physical parameters of the gait such as its perimeter and the magnitude of the connection vector fields. We use these relations to show and quantify how accurate the total Lie bracket formula is in the optimized coordinates introduced in [34].
4. A data-driven approach to obtaining the local connection for systems in granular media. This approach lays the groundwork for future studies into using the optimization framework presented in this thesis to analyze locomotion of systems where the robot-environment interaction is difficult to model from first principles.

Geometry of Locomotion

by

Suresh Ramasamy

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented August 26, 2020
Commencement June 2021

Doctor of Philosophy thesis of Suresh Ramasamy presented on August 26, 2020.

APPROVED:

Major Professor, representing Robotics and Mechanical Engineering

---

Head of the School of Mechanical, Industrial and Manufacturing Engineering

---

Dean of the Graduate School

I understand that my thesis will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my thesis to any reader upon request.

Suresh Ramasamy, Author

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to Prof. Ross Hatton for his continuous support, encouragement and guidance during my time at OSU. I would also like to thank Prof. Christine Escher for her guidance, patience and support in helping me understand concepts from differential geometry. This thesis would not have materialized without their excellent mentorship. I am extremely grateful to all my labmates, Hossein, Luke, Andrew, Jacquelin, Zach, Nathan and Quinn for all their help and friendship during my time at OSU. I thank my family and my friends for helping me stay sane during this journey. I am indebted to my parents for their unconditional love, guidance and encouragement. I am also grateful to my sister Ameeta for helping me adjust to life in America when I first moved here.

TABLE OF CONTENTS

Page

1 Introduction 1
1.1 Geometry of Optimal Gaits 4
1.2 Effect of Passive Elements on Optimal Gaits 7
1.3 Accuracy of the Total Lie Bracket Formula 9
1.4 Data-driven Connection Generation 10

2 Background 11
2.1 Geometric Locomotion Model 11
2.2 Example system dynamics models 15
2.3 Gaits 21
2.4 Finding optimal gaits 23

3 Variational Framework for Optimal Gait Identification 25
3.1 Soap-bubble Optimization of Gaits 25
3.1.1 Internal Pressure from the Lie Bracket 29
3.1.2 Surface Tension from the Distance Metric 30
3.1.3 Concentration Gradient from Parameterization 32
3.1.4 Analysis of the Purcell and serpenoid swimmers 35
3.2 Extension to Three Dimensions 38
3.2.1 Selecting a seed gait 43
3.2.2 Analysis of the four-link swimmer 44
3.2.3 Analysis of three segment piecewise continuous curvature swimmer 44
3.3 Extension to n dimensions 46
3.3.1 Selecting a seed gait 48
3.3.2 Analysis of the five-link swimmer 50
3.3.3 Comparison of different swimmers 50
3.4 Simultaneous Design and Gait Optimization 52
3.5 Conclusion 55

4 Swimmers with Elastic Tails 58
4.1 Frequency Domain Analysis 59

TABLE OF CONTENTS (Continued)

Page

4.1.1 Dynamics of the passive elastic joint . . . 61
4.1.2 Transfer function analysis . . . 63
4.2 Efficiency in the presence of passive elements . . . 68
4.2.1 Efficiency for fully actuated swimmers . . . 68
4.2.2 Efficiency for swimmers with passive elastic joints . . . 71
4.2.3 Gait parametrization for passive swimmers . . . 72
4.3 Speed-maximizing Gaits . . . 75
4.3.1 Shape gradient of the optimal input to the actuated joint. . . 77
4.3.2 Frequency gradient of the optimal input to the actuated joint 79
4.3.3 Passive Purcell and T-Link swimmers . . . 80
4.4 Energy-Efficient Gaits . . . 82
4.4.1 Shape gradient of the optimal input to the actuated joint . . 84
4.4.2 Frequency gradient of the optimal input to the actuated joint 86
4.4.3 Passive Purcell and T-Link swimmers . . . 87
4.5 Comparison with previous work . . . 90
4.6 Conclusions . . . 94

5 Magnus Expansion and Application to Gait Analysis . . . 96
5.1 Gaits and BCH formula . . . 97
5.1.1 Gaits . . . 97
5.1.2 Total Lie Bracket . . . 100
5.2 Accuracy of Total Lie Bracket . . . 106
5.2.1 Accuracy of Approximation I . . . 107
5.2.2 Accuracy of Assumption II . . . 109
5.3 Minimizing Leading Order Error . . . 110
5.4 Conclusion . . . 112

6 Data-driven Connection in Granular Media . . . 114
6.1 Previous Work . . . 115
6.2 Data-driven Connection Generation . . . 116
6.2.1 Shape and position data from encoders and markers . . . 116
6.2.2 Connection from shape and position data . . . 118
6.3 Results . . . 118

TABLE OF CONTENTS (Continued)

Page

6.4 Future Directions 119
6.4.1 Predicting displacements for systems with higher dimensional shape spaces 119
6.4.2 Inclines 122

7 Conclusion 123

Bibliography 128

Appendices 135
A Relation to sub-Riemannian Geometry 136
B Kinematics of LowRe Swimmers Obtained using RFT 139
C Comparison with Lighthill Efficiency 155
D Local Commutativity Assumption 157
E Starting Point Selection Heuristic 160

LIST OF FIGURES

Figure Page

1.1 Gaits that maximize efficiency enclose the most *curvature* of the system dynamics (measured via the curl and Lie bracket of their locomotion dynamics) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). In [59], we showed that for kinematic systems with two shape variables this process is analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble, as shown at left. More generally, the internal pressure can be thought of as being provided by the flux of the curvature passing through a surface bounded by the gait, as in the middle plot. This flux interpretation allows us to boost our geometric framework to three (right plot) or more shape variables. 5

1.2 Effect of a passive elastic joint on the shape and pacing of the optimal gait (speed-maximizing) for a Purcell swimmer. (a) The optimal gait (speed and efficiency maximizing) for the fully actuated Purcell swimmer overlaid on the curvature of the system dynamics. (b) Illustration of the Bode plot of the response of the passive elastic joint to oscillations of the active joint. This response dictates the locus of achievable gaits in the shape space for the passive swimmers. (c) The optimal gait (speed maximizing) for the passive Purcell swimmer. While the optimal gait for the fully actuated Purcell swimmer is plotted with a line of uniform thickness indicating constant power dissipation throughout the cycle, the optimal gait for the passive Purcell swimmer in (c) is plotted using a line of varying thickness, with thickness at a point corresponding to the magnitude of power required at that point of the gait. 8

2.1 Geometry and configuration variables of some of the example systems. The systems in the top row are articulated swimmers, while the bottom row consists of continuous curvature swimmers. The shape of the systems in the first columns are described by two shape variables whereas the shape of the systems in the second column are described by three shape variables. 12

LIST OF FIGURES (Continued)

Figure Page

2.2 The connection vector fields for the three-link swimmer are shown in (a) and that for the serpenoid swimmer are shown in (b). Scale in the vector fields has been chosen to emphasize structure, so scales in different components or systems are not comparable. Note that there is a qualitative similarity (modulo rotation) between the vector fields of the linked and serpenoid systems [33]. 17

2.3 Constraint Curvature Functions (rows of $D(-\mathbf{A})$) for (a) the three link swimmer and (b) the serpenoid swimmer. As with the vector fields, the CCFs exhibit a qualitative similarity (modulo a rotation) between the two sets of functions. 20

3.1 Changes in area caused by moving in the two coordinate directions in the local frame. Moving in the tangential direction $e_{\|}$ produces no change in area, as the area of the triangle given by half the product of base length and height remains the same. 26

3.2 Optimal gaits found by our algorithm overlaid onto the constraint curvature functions for their corresponding systems. Maximum-displacement gaits, which follow the zero-contours of the corresponding CCF, are indicated by dashed lines. Maximum-efficiency gaits, which are contractions of the zero-contour, are indicated by solid lines. Red regions of the CCF are positive and black regions are negative. The direction of the gait curves is matched to the sign of the region they enclose so that the displacement from each gait is positive. 33

Figure Page

3.3 Isometric embedding of the serpenoid swimmer's shape manifold in three dimensions for a drag ratio of (a) 2:1 and (b) 9:1. The surfaces in the center and right plots are constructed such that pathlengths over the surface are equal to the metric-weighted pathlengths that capture the cost of executing a gait. The red line indicates the optimal gait when the cost is calculated using the power metric $\mathcal{M}$, and the blue line indicates the optimal gait when the power metric is replaced by an identity metric in the cost calculations (treating the cost as the simple pathlength in the parameter space). Comparing the surfaces in (a) and (b) reveals that the cost manifold for the system with a drag ratio of 9 is more sharply curved than that of the system with a drag ratio of 2. This increased curvature means that the embedding of the manifold is "steeper", and that pathlength thus grows more slowly with increased parameter-radius for the 9-swimmer than the 2-swimmer, leading to a greater difference between the simple pathlength-optimum-stroke and the power-dissipation-optimum stroke for the 9-swimmer. 36

3.4 Gaits that maximize efficiency enclose the most *curvature* of the system dynamics (measured via the curl and Lie bracket of their locomotion dynamics) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). In [59], we showed that for kinematic systems with two shape variables this process is analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble, as shown at left. More generally, the internal pressure can be thought of as being provided by the flux of the curvature passing through a surface bounded by the gait, as in the middle plot. This flux interpretation allows us to boost our geometric framework to three (right plot) or more shape variables. 39

Figure Page

3.5 Changes in the area enclosed and the direction of change produced by movement along the coordinate directions in the local frame for a three-dimensional shape space. Moving in the tangential direction $e_{\parallel}$ produces no change in the area enclosed, as the area of the triangle given by half the product of base length and height remains the same. Moving in the $e_{\perp}$ direction increases the area along the $e_{\perp} \wedge e_{\parallel}$ plane and moving in the $e_{\circlearrowleft}$ direction increases the area along the $e_{\circlearrowleft} \wedge e_{\parallel}$ plane. 39

3.6 Optimal forward gaits for systems with three shape variables. (a) Four-link swimmer. The black ellipse at the center is the gait the optimizer was initialized with, and the outer red rectangle is the optimal gait found by our process. The black arrows show the vector-dual to the constraint curvature, $\overrightarrow{D(-\mathbf{A})}$ at various points. This vector-dual has an approximately constant heading over the region of the shape space explored by the optimizer and the optimal gait thus approximately evolves along a plane. (b) Three-segment piecewise-constant curvature swimmer. Top: The forward progress of the systems as they execute their optimal gait cycles. Both systems move about a tenth of a body length per cycle during these gaits, with the piecewise-constant system requiring less effort for each cycle. Insets: The time-history of the three deformation modes, illustrating a 90° phase shift between the joints on the linked system, and a smaller phase shift (more tightly-grouped peaks) between the segments of the piecewise-constant system. 41

3.7 Projection of the constraint curvature function for the 4-link swimmer onto the plane most aligned with it at its center. The blue line indicates the optimal forward gait for the 4-link swimmer. (a) The projection of the optimal gait onto this plane is a contraction of the zero contour in that plane. (b) The optimal gait for the five-link swimmer only deviates slightly from this plane, hence the zero contour of the constraint curvature function projected onto this plane provides a good seed gait for the optimizer. 45

Figure Page

3.8 The number of components of the constraint curvature in n-dimensional space is equal to the number of non-parallel two-dimensional faces on an n-cube: one for the 2-cube, or square; three for the standard 3-cube; and six for the 4-cube (or "tesseract"). (Faces on the tesseract are considered parallel if they are formed from the same pair of basis vectors, e.g., all xw faces are parallel to each other.) For higher-dimensional spaces, the number of independent two-dimensional faces continues to scale as $\frac{n(n-1)}{2}$. 47

3.9 Optimal forward gait for the five-link swimmer. The three axes represent the first three joint angle. The value of the fourth joint angle is given by the color and thickness of the line. Thin black lines indicates sections of the gait with low values of α_4 and thick red lines indicate sections of the gait with high α_4. 49

3.10 Comparison of the efficiencies of different swimmers. Red squares represent discrete link swimmers, black circles represent piecewise continuous swimmers, and blue plus signs represent serpenoid swimmers. The number next to each symbol indicates the shape variables of the system. Displacement produced by the most efficient gait for each swimmer is given by the ordinate value, and the cost of executing the gait is given by the abscissa value of the symbol representing the swimmer, and their efficiencies are the slopes of the lines connecting them to the origin. The serpenoid swimmers are the most efficient, followed by the piecewise swimmers and then the discrete link swimmers. A movie presenting a side-by-side comparison of these gaits is included in the Supplementary Material. 51

3.11 This figure shows the process of finding the optimal link length ratio and the maximum efficiency gait for the Purcell swimmer simultaneously. We seed our optimizer at the black gait (1), and the red gait (6) is the optimal gait. The inset figures provide a top view of the gaits at each iteration, and the numbers show the steps through which the gait and design variables evolve 54

LIST OF FIGURES (Continued)

Figure | Page

4.1 Comparison of the exact stroke limit cycles with the shape predicted by frequency domain analysis. Each black solid line represents the motion of the swimmer when the input to the actuated joint is a sinusoidal wave. We can see the motion converges to a limit cycle. The red dashed line represent the shape of the limit cycles predicted by frequency domain analysis presented in §4.1. The cartoons in the background show how the Purcell swimmer looks like at different points of the shape space. 60

4.2 Spring damper system whose dynamics are equivalent to the passive dynamics of the Purcell swimmer with a passive elastic joint. 63

4.3 Changes in the shape of a gait resulting from changes to amplitude and frequency of input oscillation. (a) The gait resulting from a change in amplitude of the nominal input (black) is a scaled version of the gait corresponding to the nominal input actuation (red). (b) The gait resulting from nominal input actuation. (c) The gait resulting from a change in the frequency of the nominal input (black). A change in the frequency of the nominal input leads to a change in both the amplitude and phase of the response of the passive joint. 64

4.4 Bode plot of the transfer function relating the output of the passive joint to the input of the actuated joint in the Purcell swimmer with a passive elastic tail. The inset figures show how the periodic orbits corresponding to gain and phase at certain frequencies look like in the shape space. We can see that actuation at very low frequencies leads to gaits that enclose very little surface area due to the amplitude of the passive joint being low, while at very high frequencies the surface area enclosed by the gait is low due to the almost 180 degree phase shift between the oscillations of the actuated and passive joints. 65

4.5 A flowchart describing how the parameters describing our input to the actuated joint affect the two optimality criteria in the case of drag-dominated swimmers with a passive joint. 72

Figure Page

4.6 Gaits that maximize speed along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated Purcell swimmers at different power levels. 81

4.7 Gaits that maximize speed along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated T-link swimmers at different power levels. 81

4.8 Gaits that maximize efficiency along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated Purcell swimmers moving forward at the same speed as the passive swimmer. Note that for the passive Purcell swimmer, the optimizer stops because reducing the frequency further or making the gait smaller does not provide any meaningful increase in efficiency. This observation is line with the results in [56]. 85

Figure Page

4.9 Gaits that maximize efficiency along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated T-link swimmers moving forward at the same speed as the passive swimmer. 85

4.10 In all the subfigures, the solid red lines and the red circles show the speeds and efficiencies predicted by numerical simulation and integral of CCF respectively. The solid black lines show the speeds and efficiencies predicted by the constant-CCF assumption used in [56] in the link-attached coordinates. The first column of figures illustrate the speed of the (a) Purcell swimmer as a function of actuation frequency when the input to the controlled joint is a sinusoidal oscillation of unit amplitude and when the input to the controlled joint has the optimal shape and (b) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §4.3. In (a), the grey line shows the speed predicted by the constant-CCF assumption for a sinusoidal input. The small amplitude perturbation analysis in [56] predicts speed to be a monotonically increasing function of frequency for all inputs. The speed is a monotonically increasing function of frequency for the sinusoidal input. However, it is not monotonically increasing for the optimal input contrary to the prediction from the small amplitude perturbation analysis in [56]. The second column of figures illustrate the efficiency of the (c) Purcell swimmer and (d) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §4.4. Figures (e)-(h) illustrate the speed and efficiency of the Purcell swimmer and T-link swimmer as a function of the magnitude of optimally shaped input actuation obtained from §4.4. 89

LIST OF FIGURES (Continued)

Figure Page

5.1 A square image family of gaits with side length, $L = 1$, overlaid on the rotational connection vector field for the differential drive car in the original and optimal coordinates are shown in (a) and (c) respectively. Similar rotational connection vector field plots for the kinematic snake are shown in (e) and (g) respectively. The black boxes with numbers 1 and 3 indicate the starting points which minimize the leading order error and the blue boxes with numbers 2 and 4 indicate the starting points which maximize the leading order error. The displacement produced by executing the square gait from these starting points are shown in the figures adjacent to the connection vector field plots. The red circles in (b),(d),(f) and (g) show where the systems starts out from and the red star shows the displacement predicted by the total Lie bracket formula. The magnitude of leading order error E depends on the starting point of the gait. The red dashed lines represent the range of leading order error E. Starting the gait from locations denoted by black boxes produce the least amount of error, while starting from blue boxes maximizes the error as expected. The range of leading order error E is also smaller in the optimal coordinates than in the original coordinate used to describe the systems. 111

6.1 Motion of the SEA snake from all datasets imposed on a single figure. These motions were used to experimentally determine the local connection $\mathbf{A}$ for the SEA snake. 117

6.2 Comparison of the connection vector fields obtained through RFT for the 9 link viscous swimmer with the connection vector fields obtained through data driven analysis for the SEA snake. The connection vector fields show many qualitative similarities. 120

6.3 A comparison of the predicted x-directional body velocity by our data-driven vectorfield (black dotted line) and the observed x-directional body velocity when executing spiral gaits. 121

6.4 A comparison of the predicted y-directional body velocity by our data-driven vectorfield (black dotted line) and the observed y-directional body velocity when executing spiral gaits. 121

LIST OF FIGURES (Continued)

Figure Page

6.5 A comparison of the predicted rotational body velocity by our data-driven vectorfield (black dotted line) and the observed rotational body velocity when executing spiral gaits. 121

Chapter 1: Introduction

Many animals and robots move through the world by coupling cyclical changes in shape called gaits to an interaction with the environment. The major goal of this thesis is to build a framework that relates efficiency of gaits to geometric quantities associated with the gait such as its perimeter and the area enclosed by the gait. We use this framework to identify optimal gaits for swimmers in drag-dominated environments and understand the impact of passive elements like springs on the efficiency of swimmers.

Robot locomotion is a significant area of research in robotics. The design of locomotors for robots draws heavily from locomotion methods employed successfully by different animals. This is particularly true for robots operating in harsh environments with either heterogenous ground or deformable substrates like sand, mud, and soil. Hence the diversity in locomotion techniques employed by robots mirrors the diversity observed in animal locomotion. A major stumbling block while studying robot locomotion is that most strong results are limited by a set of assumptions (such as the quasi-static assumption) or are applicable to only a particular robot morphology (such as bipeds or quadrupeds). The lack of a universal framework impedes comparison of different approaches to locomotion and hinders discovering fundamental principles that apply to all locomoting systems.

One approach to developing a unifying framework to analyze, design and control

robotic locomotion is to determine the normal form or a small number of normal forms of the equations of motion (EOM) that can capture the mechanics of a large class of locomoting systems. Here a normal form refers to a set of equations exhibiting certain underlying structure. Systems whose EOM can be put into a common normal form share the underlying structure of the normal form and only the numerical values of the terms constituting the equations change between individual systems. This approach has the potential to serve as the basis for creating useful tools for analyzing locomotion analogous to how the normal form $\dot{x} = Ax + Bu$ in linear control theory helped build a substantial body of results that help analyze and control a wide range of electromechanical systems.

A normal form that describes a broad class of systems allows for evaluation and comparison of efficiency of locomotion across the class in a way that is independent of how the forces were generated in individual cases. These comparisons help uncover fundamental principles underlying locomotion for systems in the class and provides insights greater than those obtained from the analysis of individual systems similar to how the Spring-Loaded Inverted Pendulum (SLIP) model helped the study of hopping, jumping and running gaits across different robots.

With this goal in mind, a normal form of the equations that describe mechanics of locomotion when the interaction with the environment was presented in [55]. The purely kinematic version of the normal form in [55] captures the structure of the underlying motion for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [64, 68], fluid interactions at the extremes of low [4, 33] and high [33, 39, 46] Reynolds numbers, and explicit

nonholonomic constraints such as passive wheels [11, 52, 53, 64].

The normal form in [55] helped show that many modes of locomotion can be modelled using mathematical objects called connections from principle bundle theory. This differential geometric insight helped develop an intuitive understanding of how systems locomote and gave us tools that used the curvature of system dynamics (encoded in their Lie brackets) to identify classes of shape oscillations that resulted in desirable net displacements.

An enduring goal [4, 40, 46, 51, 58, 64, 68] of the geometric mechanics community has been to go beyond simply identifying the direction in which small shape oscillations will propel the system, and to use these curvature methods to identify and geometrically describe optimal gaits that generate the most displacement, either per gait cycle or per energy expended. A key obstacle to reaching this goal was the noncommutativity in the equations of motion for most interesting locomoting systems. Because translations and rotations do not commute with each other, summing up Lie brackets over a finite region provides only an approximation of the displacement induced by a gait, and the error in this approximation increases with the amplitude of the oscillation. Historically (e.g., [46]), this error was regarded as growing too quickly to provide any meaningful information about the optimal gait cycles. In [29–31, 34], however, the authors introduced a choice of coordinates that converts much of the system's noncommutativity into nonconservativity, which *is* amenable to finite-scale integration. Working in these new coordinates has provided us insights into the optimal gaits of swimmers and crawlers, in low- and high-viscosity fluids [28, 32, 33], and in granular media [27]. Further, these in-

sights can be used to easily "hand-draw" optimized gaits that would otherwise only have been found by a long numerical search of a high-dimensional space of trajectory parameters [27].

1.1 Geometry of Optimal Gaits

In this thesis, we consolidate and extend the geometric insights from previous work into a set of geometric principles that govern the geometry of optimally-efficient gaits for drag-dominated systems with an arbitrary number of shape variables. We formally encode these principles in a set of geometric expressions that together make up the gradient of the gait efficiency with respect to variations in the gait trajectory. This gradient can be used directly in a gradient-descent solver to find optimal gaits (as in the examples we provide in the text), but more fundamentally, *the gradient geometrically describes the dynamics underlying any other gait optimization algorithm applied to the system.*

The core elements of the gait gradient are:

1. A gradient ascent/descent component that pushes the cycle to enclose a large sign-definite region of the constraint curvature, maximizing the net displacement generated by the gait;
2. A cost component based on a Riemannian metric that limits the growth of the gait cycle; and
3. A perimeter-balancing component that evenly spaces the points along the trajectory, stabilizing the solution and providing an efficiency-optimal pa-

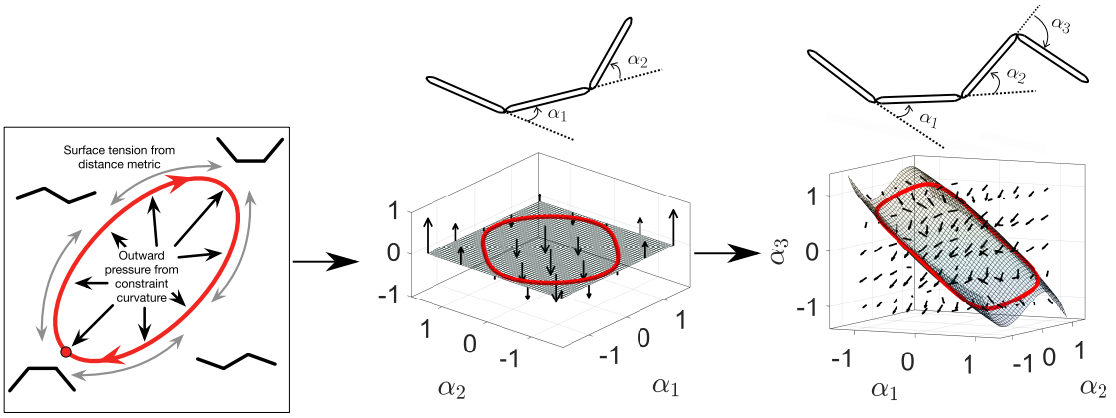

Figure 1.1: Gaits that maximize efficiency enclose the most curvature *of the system dynamics (measured via the curl and Lie bracket of their locomotion dynamics) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). In [59], we showed that for kinematic systems with two shape variables this process is analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble, as shown at left. More generally, the internal pressure can be thought of as being provided by the flux of the curvature passing through a surface bounded by the gait, as in the middle plot. This flux interpretation allows us to boost our geometric framework to three (right plot) or more shape variables.*

rameterization of the resulting motion.

As illustrated in Fig. 1.1, the dynamics of this optimizer are similar to the dynamics of a soap bubble, with the Lie bracket providing an internal pressure which causes the gait cycle to expand, the metric-weighted pathlength providing a surface tension that constrains the growth of the gait, and the perimeter-balancing term providing a concentration gradient that evenly distributes waypoints along the gait.

With this basic soap-bubble framework in place for systems with two shape variables, we next extend our formulation to three dimensions by recasting some

of the geometric terms, as we initially described in [60]. In particular:

1. We show how the constraint-curvature "area enclosure" generalizes to a flux integral through an oriented surface (more formally, an integral over a differential two-form), as illustrated in Fig. 1.1;
2. We investigate how optimal gaits "bend" to exploit this orientation and maximize the flux passing through the surface.

Finally, we consider the structure of the "flux integral" term for systems with n shape variables. Beyond three dimensions, the convenient vector-flux analogy from two and three dimensions no longer holds, but fully-adopting the formalism of differential-geometric two-forms allows us to extend our method to these spaces.

As a demonstration of this approach, we use it to identify optimal gaits for a set of example systems moving in viscosity-dominated (low Reynolds number) environments. These systems include Purcell's three-link swimmer [57] (a standard benchmark for locomotion analysis), four- and five-link swimmers, and several continuous-curvature extensions of Purcell's swimmer with different classes and numbers of shape modes [35]. We identify optimal gaits for these systems in both the forward and turning directions, which match those found previously via raw parameter optimization in works such as [28, 66]. The optimal gaits we find for these systems highlight the benefits of a continuous backbone. We then present how this framework can also be used for design optimization by optimizing link length ratio and joint kinematics simultaneously.

1.2 Effect of Passive Elements on Optimal Gaits

We then extend this framework to identify optimal gaits for swimmers with passive elastic joints. Drag-dominated swimmers with passive elements were studied in [15, 38, 42, 49, 56]. Of these works, [42, 56] and [49] are most relevant to us, as they discuss the motion of swimmers with a harmonically-driven active joint and a passive joint. The analysis in [56] is particularly relevant, where using perturbation expansion, explicit expressions for leading order solutions were derived for harmonic input oscillations, and the optimal swimmer geometry was obtained for the Purcell swimmer. In this work, unlike in [56], we do not restrict our input to simple harmonic oscillations, and use a higher order representation of the system dynamics.

The ways in which the dynamics of swimmers with passive elastic joints differ from the dynamics of fully actuated swimmers are:

1. The motion of the actuated and passive joints are coupled in the passive swimmers, and hence passive swimmers can execute only some of the gait kinematics the fully actuated swimmers can execute.
2. Due to this coupling, there is also a unique pacing associated with each gait that the passive swimmer can execute, so the swimmer cannot minimize the pacing cost separately from the kinematic cost.

As illustrated in Fig. 1.2, we address the two problems introduced by the presence of passive elastic joints: we first use frequency domain analysis to analytically

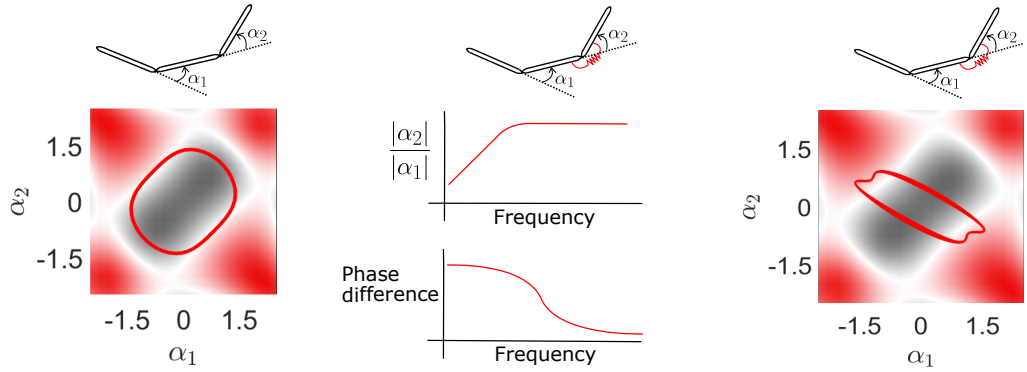

(a) Optimal gait for the fully actuated Purcell Swimmer

(b) Frequency response of the passive joint to oscillations of the active joint

(c) Optimal gait for the passive Purcell swimmer

Figure 1.2: Effect of a passive elastic joint on the shape and pacing of the optimal gait (speed-maximizing) for a Purcell swimmer. (a) The optimal gait (speed and efficiency maximizing) for the fully actuated Purcell swimmer overlaid on the curvature of the system dynamics. (b) Illustration of the Bode plot of the response of the passive elastic joint to oscillations of the active joint. This response dictates the locus of achievable gaits in the shape space for the passive swimmers. (c) The optimal gait (speed maximizing) for the passive Purcell swimmer. While the optimal gait for the fully actuated Purcell swimmer is plotted with a line of uniform thickness indicating constant power dissipation throughout the cycle, the optimal gait for the passive Purcell swimmer in (c) is plotted using a line of varying thickness, with thickness at a point corresponding to the magnitude of power required at that point of the gait.

approximate the motion of the unactuated joint in response to the motion of the actuated joint. We then combine this frequency-space analysis with elements of the geometric framework introduced in [61] to construct a gradient-descent algorithm that identifies optimal gaits and the pacing associated with these gaits for passive swimmers. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

1.3 Accuracy of the Total Lie Bracket Formula

The total Lie bracket formula obtained by restricting the BCH expansion to the first few terms forms the backbone of the optimization framework presented in this thesis. Hence, it is crucial that the truncation be a good approximation of the full BCH expansion. The aim of chapter 5 is to understand the factors affecting the accuracy of the total Lie bracket formula. We show that the major factors affecting the accuracy are the magnitude of the connection vector field and the size of the gait. Since the optimal coordinate choice minimizes the magnitude of the connection vector field, it improves the accuracy of the truncated Magnus series and hence the accuracy of the total Lie bracket formula.

1.4 Data-driven Connection Generation

Another key goal of our geometric study of locomotion is to make this approach feasible for robots whose dynamics are difficult to model from first principles, as in the granular systems of [16, 27]. In chapter 6 we present a process for obtaining a data-driven connection for a robot operating in granular media. The dimensionality extensions and explicit optimality criteria discussed in this thesis will broaden the classes of systems that we can consider in the empirical-geometric work. Additionally, the geometric structure we discuss here forms the foundation of a geometric-Floquet approach to online data-driven locomotion analysis and optimization on very-high dimensional systems [10].

Chapter 2: Background

In this chapter we will review how the dynamics of our example systems are obtained and some of the geometric insights from our previous work. Our treatment below is condensed from a series of papers we have written for the robotics community [31, 33, 35], and at a deeper mathematical level, in [34].

For the purposes of this work, our focus is on the geometric structure of the system dynamics. Accordingly, we work with the components of these dynamics at a relatively high level of abstraction in the equations, and present their instantiation for specific systems graphically rather than as algebraic expressions (which would run to several pages of trigonometric terms if expanded, even for the three-link swimmer). For worked examples of the construction of the dynamics of the three-link swimmer, see [33], [35] and appendix B.

2.1 Geometric Locomotion Model

When analyzing a locomoting 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

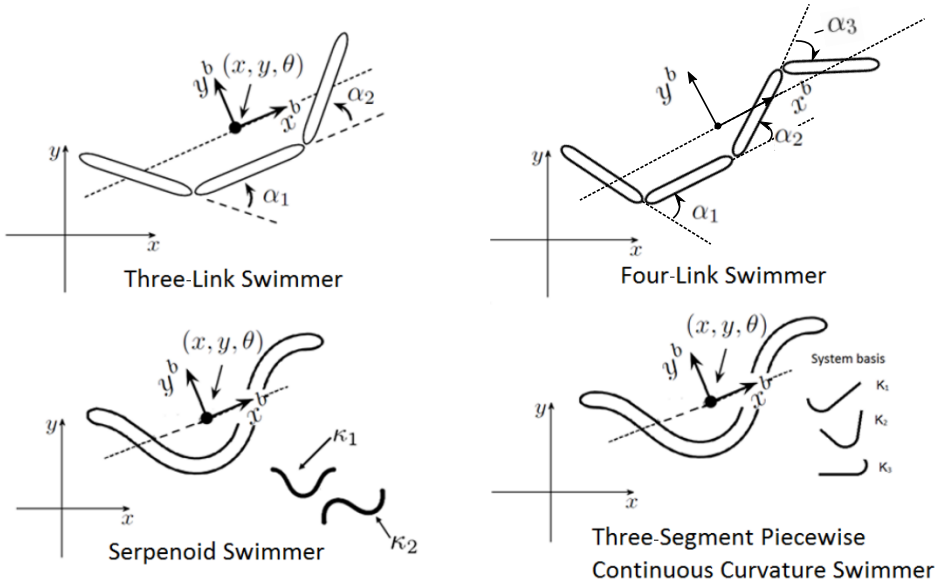

Figure 2.1: Geometry and configuration variables of some of the example systems. The systems in the top row are articulated swimmers, while the bottom row consists of continuous curvature swimmers. The shape of the systems in the first columns are described by two shape variables whereas the shape of the systems in the second column are described by three shape variables.

it.[1] For example, the positions of both the articulated and continuous-curvature swimmers in Fig. 2.1 are the locations and orientations of their centroids and mean orientation lines, $g = (x, y, \theta) \in SE(2)$. The shape of the articulated swimmers are parameterized by their joint angles, $r = (\alpha_1, \alpha_2)$ for the three link swimmer and $r = (\alpha_1, \alpha_2, \alpha_3)$ for the four-link swimmer. The shape of the continuous curvature swimmers can be described by a set of modal amplitudes multiplied by the curvature modes. In the serpenoid and piecewise-continuous systems, the shape parameters α are weighting functions on curvature modes κ defined along the body, as discussed in [35].

A useful model for locomotion in *kinematic* regimes where no gliding can occur,[2] and which we employ in this work, is that at each shape, there exists a linear relationship between changes in the system's shape and changes in its position,

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

in which $\mathring{g} = g^{-1}\dot{g}$ is the body velocity of the system (i.e., $\dot{g}$ expressed in the system's local coordinates), and the *local connection* $\mathbf{A}$ acts like the Jacobian of a robotic manipulator, mapping from joint velocities to the body velocity they produce by pushing the system against its environment. Each row of $\mathbf{A}$ can be regarded as a body-coordinates local gradient of one position component with

[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* [34].

[2] This kinematic condition has been demonstrated for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [64, 68], non-holonomic constraints such as passive wheels [11, 52, 53, 64], and fluid interactions at the extremes of low [4, 33] and high [33, 39, 46] Reynolds numbers.

respect to the system shape. If we plot the rows of $\mathbf{A}$ as arrow fields, as in Fig. 2.2, this means that moving in the direction of the arrows moves the system positively in the corresponding body direction, and moving perpendicular to the arrows results in no motion in that direction [31, 34].

In a drag-dominated environment, the effort required to change shape can be modeled as the pathlength s of the trajectory through the shape space, weighted by a Riemannian metric $\mathcal{M}$ as

$$ds^2 = dr^T \mathcal{M}\ dr. \tag{2.2}$$

Inline with our previous work in [32], we take the metric tensor $\mathcal{M}$ as the mapping from shape velocity to power dissipated into the surrounding medium,

$$P = \dot{r}\mathcal{M}\dot{r}, \tag{2.3}$$

which is encoded by the same matrix as the mapping from joint velocities to torques on the joints

$$\tau = \mathcal{M}\dot{r}, \tag{2.4}$$

and so can be readily calculated from the first-principles physics of the system.

As discussed in [3, 7, 35], the length of a path under this metric describes the time it takes for a system dominated by viscous drag to follow the path at unit power (or, equivalently, the square root of the power required to execute it in unit time). Because moving with constant power is the least-costly pacing with

which to execute a motion under viscous drag [7], this pathlength thus provides a geometric cost for the best-case execution of the kinematics in a gait cycle.

Any pacing other than constant-power will make the trajectory take longer for a given average power (or increase the average power required to complete the motion in a fixed time). The additional cost for a non-optimal pacing can be represented by squaring the difference between the average and instantaneous rates at which the gait is being followed (measured as s per time), and then integrating over the time during which the gait is being executed,

$$\sigma = \int_0^{\tau_{total}} \left(\frac{s_{\text{total}}}{\tau_{\text{total}}} - \frac{d}{d\tau}(s(\tau))\Big|_{\tau=t} \right)^2 dt, \tag{2.5}$$

where τ_{total} is the time period of the gait, s_{total} is the length of the gait under the metric $\mathcal{M}$, and s is distance traveled along the gait as a function of time corresponding to the given pacing. If the gait is proceeding at constant power, $\frac{s_{\text{total}}}{\tau_{\text{total}}}$ is equal to the rate at which s changes with time, so σ measures the extent to which the pacing lags and leads the optimal pacing.

2.2 Example system dynamics models

In this work, as in [33, 35], we generate the dynamics for our example systems from a resistive force model, in which each element of the body is subject to normal and tangential drag forces proportional to their velocities in those (local) directions.[3]

[3]This model is most widely associated with swimmers at low Reynolds numbers (e.g. [66]), but can also be regarded as an informative general model for systems that experience more lateral drag than longitudinal drag (e.g. [27]). Our choice of resistive force here also does not preclude

The normal drag coefficient is larger than the tangential component (here, by a factor of 2 : 1), corresponding to the general principle that it is harder to move a slender object in a fluid or on a surface crosswise than it is to move it along its length. We then impose a quasi-static equilibrium condition that the net drag force and moment on the system is zero at all times (treating the system as heavily overdamped, with acceleration forces much smaller than drag forces); because the drag forces are not isotropic, the system can use the angle-of-attack of its body surfaces to generate net motion.

Together, these conditions impose a *Pfaffian constraint*[4] on the system's generalized velocity,

$$\begin{bmatrix} F_x^b \\ F_y^b \\ F_\theta^b \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \omega(r) \begin{bmatrix} \mathring{g} \\ \dot{r} \end{bmatrix}, \tag{2.6}$$

in which the matrix ω that maps the velocities to the net forces on the body frame is a function of the system's internal kinematics and depends only on the shape r.[5]

By separating ω into two sub-blocks, $\omega = [\omega_g^{3\times 3}, \omega_r^{3\times n}]$, it is straightforward to

the use of more detailed physical models (e.g., [24]) to construct the local connection $\mathbf{A}$.

[4] A constraint that the allowable velocities are orthogonal to a set of locally-linear constraints, i.e., that they are in the nullspace of a constraint matrix ω.

[5] The expressions for the dynamics are unwieldy (running to several pages of trigonometric terms in even the simplest cases) so we do not write them out in full here. See [33] for a more detailed treatment of (2.6)–(2.7) in the case of the three-link swimmer, and [35] for how we build the metric tensor for that system.

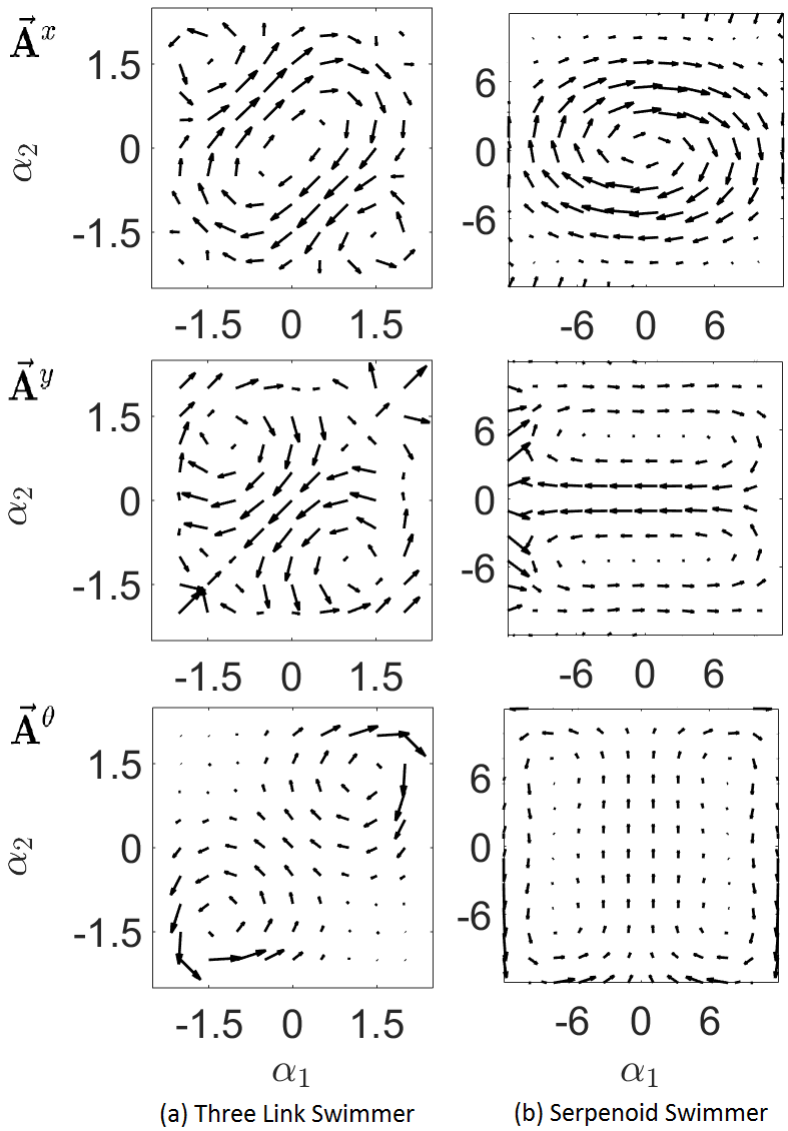

Figure 2.2: The connection vector fields for the three-link swimmer are shown in (a) and that for the serpenoid swimmer are shown in (b). Scale in the vector fields has been chosen to emphasize structure, so scales in different components or systems are not comparable. Note that there is a qualitative similarity (modulo rotation) between the vector fields of the linked and serpenoid systems [33].

rearrange (2.6) into

$$\mathring{g} = -(\omega_g^{-1}\omega_r)\dot{r}, \tag{2.7}$$

revealing the local connection as $\mathbf{A} = \omega_g^{-1}\omega_r$. For a more detailed treatment this process, we refer the reader to [33]. Once $\mathbf{A}$ has been found, it can be used to calculate a Riemannian metric $\mathcal{M}$ over the shape space as

$$\mathcal{M}(r) = \int_{\text{body}} J^T(r,\ell)\ c\ J(r,\ell)\ d\ell, \tag{2.8}$$

where $J(r,\ell)$ is the Jacobian from shape velocity to the local velocity of each section of the body (which incorporates both $\mathbf{A}$ and the system's internal kinematics), and c is the matrix of drag coefficients, which acts as a local metric for the motion of each link. For the systems considered in the paper it is given explicitly by,

$$c = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{2.9}$$

indicating that for any infinitesimal element of a link, the resistance to lateral motion is twice the resistance to longitudinal motion. See [36] for more details on how we build the metric tensor.

As discussed in our previous work [33, 34], the metric $\mathcal{M}$ encodes a quadratic relationship between the shape velocities and power dissipated into the surroundings, given by

$$P = \dot{r}^T\mathcal{M}(r)\dot{r}, \tag{2.10}$$

as well as the mapping from joint velocities to torques on the joints,

$$\tau = \mathcal{M}(r)\dot{r} \tag{2.11}$$

Details of the calculations to generate the local connection $\mathbf{A}$ and the Riemmanian metric $\mathcal{M}$ for our example systems are provided in Appendix B.

We apply this physics model to several example geometries, which are illustrated in Fig. 2.1: Purcell's three-link swimmer [57], a four-link swimmer, a *serpenoid swimmer* [37] and piecewise-constant curvature swimmers with two, three, and four segments. The three-link swimmer is a useful and widely-adopted [4, 8, 18, 23, 33, 46, 66] minimal example for locomotion, because its two degrees of freedom can be easily visualized. The serpenoid swimmer, whose shape is defined by the amplitude of sine and cosine curvature modes [32], provides an example of a two-DoF system that has been shown to closely model how snakes and other animals use undulatory locomotion to move through the world [37]. The piecewise-constant curvature geometries are intermediaries between the discrete joints of the linked systems and the smooth traveling waves of the serpenoid systems, and correspond to system morphologies that are achievable via simple soft robotic actuators [12].

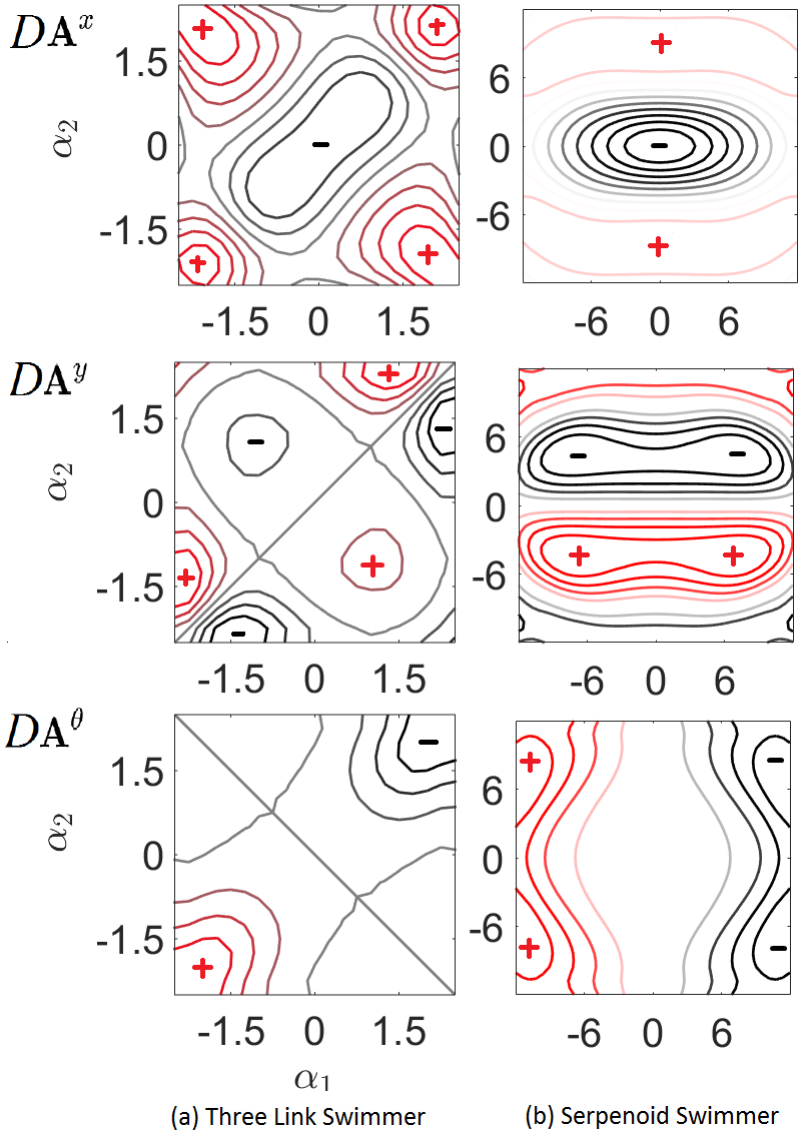

Figure 2.3: Constraint Curvature Functions (rows of $D(-\mathbf{A})$) for (a) the three link swimmer and (b) the serpenoid swimmer. As with the vector fields, the CCFs exhibit a qualitative similarity (modulo a rotation) between the two sets of functions.

2.3 Gaits

Because the shape space of locomoting systems is typically bounded (e.g., because of joint limits or other restrictions on bending the body), these systems often move via gaits: cyclic changes in shape that produce characteristic net displacements. Several efforts in the geometric mechanics community [4, 46, 50, 52, 53, 64, 68] (including our own [32, 33]), have aimed to use *curvature* of the system constraints (a measure of how "non-canceling" the system dynamics are) to understand which gaits produce useful displacements.

The core principle in these works is that because the net displacement g_ϕ over a gait cycle ϕ is the line integral of (2.1) along ϕ, the displacement can be approximated[6] by an area integral of the curvature $D(\text{–}\mathbf{A})$ of the local connection (its total Lie bracket [34]) over a surface ϕ_a bounded by the cycle:

$$g_\phi = \oint_\phi -g\mathbf{A}(r) \tag{2.12}$$

$$\approx \iint_{\phi_a} \underbrace{-\mathbf{dA} + \sum\left[\mathbf{A}_i, \mathbf{A}_{j>i}\right]}_{D(\text{–}\mathbf{A}) \text{ (total Lie bracket)}}, \tag{2.13}$$

where $\mathbf{dA}$, the exterior derivative of the local connection (its generalized row-wise curl), measures how changes in $\mathbf{A}$ across the shape space prevent the net induced motions from canceling out over a cycle, and the local Lie bracket $\sum\left[\mathbf{A}_i, \mathbf{A}_{j>i}\right]$

[6] This approximation (a generalized form of Stokes' theorem) is a truncation of the Baker-Campbell-Hausdorf series for path-ordered exponentiation on a noncommutative group, and closely related to the Magnus expansion [45, 58]. The accuracy of this approximation depends on the body frame chosen for the system, whose selection we discuss in [33, 34]. In presenting this approximation, we also elide some details of exponential coordinates on Lie groups, which are also discussed in [34].

measures how translations and rotations in the induced motions couple into "parallel parking" effects that contribute to the net displacement.

For systems with two shape variables, the exterior derivative term evaluates as

$$\mathbf{dA} = \left(\frac{\partial \mathbf{A}_2}{\partial r_1} - \frac{\partial \mathbf{A}_1}{\partial r_2} \right) d\alpha_1 \wedge d\alpha_2 \tag{2.14}$$

and the local Lie bracket term for planar translation and rotation evaluates as

$$\left[\mathbf{A}_1, \mathbf{A}_2\right] = \begin{bmatrix} \mathbf{A}_1^y \mathbf{A}_2^\theta - \mathbf{A}_2^y \mathbf{A}_1^\theta \\ \mathbf{A}_2^x \mathbf{A}_1^\theta - \mathbf{A}_1^x \mathbf{A}_2^\theta \\ 0 \end{bmatrix} d\alpha_1 \wedge d\alpha_2. \tag{2.15}$$

In both cases, the wedge product $d\alpha_1 \wedge d\alpha_2$ indicates the oriented differential area basis in the shape space.

Plotting these curvature terms as scalar functions over the shape space (as in Fig. 2.3) reveals the effect of gaits' geometry on the motion they induce: Gaits that produce large net displacements in a given (x, y, θ) direction are located in strongly sign-definite regions of the corresponding $D(-\mathbf{A})$. For example, x-translation gaits encircle the center of the shape space for both the three-link and serpenoid systems, whereas y-translations or θ-rotations are produced by cycles in the corners or edges of the shape space respectively [15].

2.4 Finding optimal gaits

Optimal gait design has a long history of research in the physics, mathematics, and engineering communities, as part of the broader field of optimal control [14, 54]. For systems of the classes we consider here, notable contributions include those of Purcell, who introduced the three-link swimmer as a minimal template for understanding locomotion, a series of works [8, 18, 23, 65, 66] aimed at numerically optimizing the stroke pattern, the observation in [7] that the optimal pacing for the gait keeps the power dissipation constant over the cycle, the recognition that the gait optimization problem can be formulated as a variational optimization problem with a perimeter-length cost [3, 4] and work on finding optimal gaits through two-dimensional slices of higher-dimensional shape spaces [26]. Reasonably efficient gaits are presented in [71] as natural oscillations of a locomoting system, defined as the free periodic response obtained when the damping effect is partially reduced so that the system becomes marginally stable. Optimally swimming between two points can also be considered as swimming along a subriemmanian geodesic [48], and [2] presents a shooting based method to identify these geodesics. [1, 9] present numerical schemes to find these geodesics for swimming mechanisms.

Our geometric view of the system dynamics allows us to make two strong statements about the nature of optimal gaits:

1. The maximum-displacement gaits (the "longest strides" that the systems can take) follow the zero-contours of the constraint curvature functions, completely enclosing a sign definite region.

2. The most efficient gaits (where we define efficiency η as the ratio of the displacement produced by the gait to the cost of executing the gait, $\frac{g_\phi}{s}$), are contractions/straightenings of these zero-contours as shown in Fig. 3.2, and correspond to the systems' "comfortable strides": by giving up the low-yield regions at the edges of the sign-definite regions (or crossing slightly into opposite-sign regions) the system travels a shorter distance in each cycle, but the consequently shorter perimeter length means the system can repeat the cycles more quickly at a given level of power consumption.[7]

Our definition of efficiency is equivalent to the inverse of the mechanical cost of transport used in [3] (where we prefer to work in terms of average velocity achieved at a given power level instead of power required to maintain a desired average speed). Gaits that optimize our criterion also optimize Lighthill's efficiency, which compares the power dissipated while executing the gait to the power dissipated in rigidly translating the swimmer through the fluid.

[7]This definition of efficiency is invariant with respect to time reparameterizations of the gait curve, in that it assumes that the gait curve will be reparameterized to follow at a best-case (constant-power) pacing at execution time, and in that the relative efficiency of two gaits remains the same under bulk rescaling of time. The efficiency is not, however, dimensionless with respect to system parameters such length-scale or drag coefficients, which both affect the values in the metric tensor $\mathcal{M}$ as discussed in [35].

Chapter 3: Variational Framework for Optimal Gait Identification

In this chapter, we present our framework for identifying optimal gaits for fully actuated drag-dominated swimmers. In §3.1 we present our variational optimizer formulation for systems with two shape variables. In §3.2 we extend this formulation to systems with three shape variables. In §3.3 we extend this formulation to systems with more than three shape variables. In §3.4 we examine how our variational formulation can be used to concurrently optimize a system's physical structure and the gaits it executes. Appendix A relates the work in this paper to prior work on these systems that made use of subriemannian geometry, Appendix B provides more details on computing the system dynamics for our examples

3.1 Soap-bubble Optimization of Gaits

In this section, we encode the geometric principles described in §2.4 into a gait optimization algorithm for fully actuated drag dominated swimmers and show that for systems with two shape variables, the process of identifying the optimal gait is analogous to how a soap bubble obtains its shape.

We start from the basic variational principle that functions reach their extrema when their derivatives go to zero. Given a gait parameterization p, maximum-displacement cycles therefore satisfy the condition that the gradient of net dis-

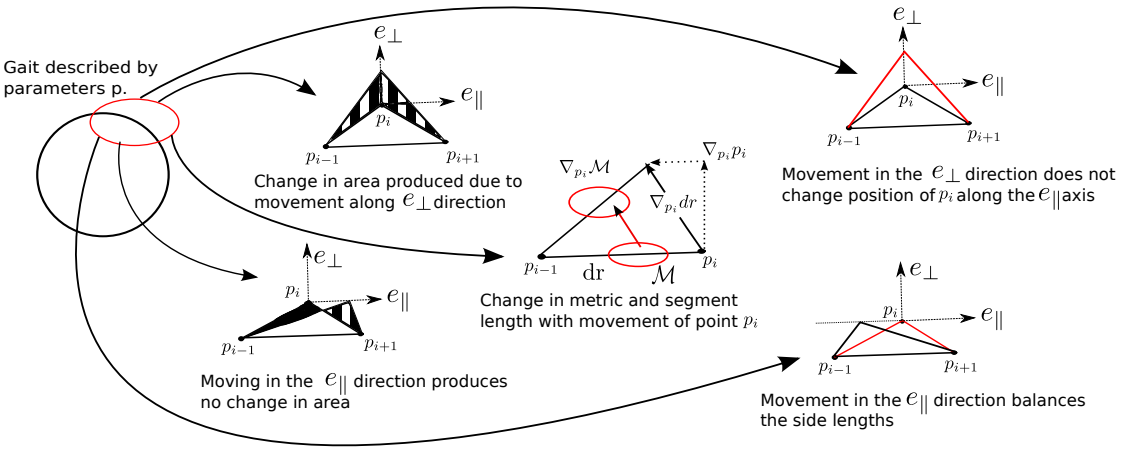

Figure 3.1: Changes in area caused by moving in the two coordinate directions in the local frame. Moving in the tangential direction $e_\parallel$ produces no change in area, as the area of the triangle given by half the product of base length and height remains the same.

placement with respect to the parameters is zero,

$$\nabla_p g_\phi = \mathbf{0}. \tag{3.1}$$

Maximum-efficiency cycles likewise satisfy the condition that the gradient of the efficiency ratio is zero,

$$\nabla_p \frac{g_\phi}{s} = \frac{1}{s}\nabla_p g_\phi - \frac{g_\phi}{s^2}\nabla_p s = \mathbf{0}. \tag{3.2}$$

where the efficiency, as described in Section §2.4, is taken as the displacement g_ϕ normalized by the pathlength-effort required to execute the cycle (s as calculated in (2.2)).

For suitable seed values p_0, solutions to (3.1) and (3.2) can therefore be reached

by finding the respective equilibria of the dynamical systems

$$\dot{p} = \nabla_p g_\phi \quad \text{and} \quad \dot{p} = \nabla_p \frac{g_\phi}{s}. \tag{3.3}$$

The stable equilibria of the right-hand equation in (3.3) are gaits in the same "image families" as the system's optimally-efficient gaits (i.e., they follow the same curve as the optimal gait, but not necessarily at the same pacing). To construct the optimal gait, we can either optimize via (3.3) and then choose a constant-metric-speed parameterization, such that the pacing penalty σ from (4.20) goes to zero, or directly include $\nabla_p \sigma$ in our optimizer.[1]

Combining the gradient of the pacing term with the gradient of the image-optimizer places the maximum-efficiency gait as the equilibrium of

$$\dot{p} = \nabla_p g_\phi - \frac{g_\phi}{s} \nabla_p s + \nabla_p \sigma \tag{3.4}$$

(from which we have factored out a coefficient of $\frac{1}{s}$ from (3.2)). As illustrated in Fig. 3.4, this differential equation is directly analogous to the equations governing the shape of a soap bubble: $\nabla_p g_\phi$ takes the Lie bracket as an "internal pressure" seeking to expand the gait cycle to fully encircle a sign-definite region, $\nabla_p s$ is the "surface tension" that constrains the growth of the bubble, and $\nabla_p \sigma$ is the

[1] Including $\nabla_p \sigma$ in the optimizer works best for parameterizations in which $\nabla_p \sigma$ is orthogonal to $\nabla_p \frac{g_\phi}{s}$, such as waypoint based direct transcriptions. For other parameterizations, e.g., Fourier series, the gradients may not be orthogonal and a two-step procedure of optimizing the image then the pacing will produce better results. For waypoint-based parameterizations, the $\nabla_p \sigma$ term has a secondary benefit of helping to stabilize the optimizer by maintaining an even spacing of points and thereby preventing the formation singularities in the curve.

"concentration gradient" that spreads the soap over the bubble's surface. In the following subsections, we explore each of the terms in (4.33), discussing both their fundamental geometric definitions and how they can be implemented in a direct-transcription solver.

In our examples, we parametrize the gait as a sequence of waypoints p_i such that the gait parameters p_i explicitly define the location of the discretization points. As illustrated in Fig. 3.1, each waypoint p_i forms a triangle with its neighboring points and we can define a local tangent direction $e_{\|}$ as

$$p_{i+1} - p_{i-1} = \ell\, e_{\|} \tag{3.5}$$

and a local normal direction $e_{\perp}$ orthogonal to $e_{\|}$.

We select this direct-transcription parameterization because it facilitates visualizing the workings of our optimizer (and thus the dynamics governing any other optimization applied to this problem). Additionally, it allows us to illustrate simultaneous optimization of the gait path and its pacing. We could also parametrize the gait using a Fourier series or Legendre polynomials. In this case, the pacing optimization would have to be done after the image of the optimal gait has been found because finding an optimal pacing can no longer be formulated as a process orthogonal to the gradient descent process for finding the image of the optimal gait.

3.1.1 Internal Pressure from the Lie Bracket

The first term in (4.33), $\nabla_p g_\phi$, guides the gait towards maximum-displacement cycles. By sustituting the approximation from (5.5) into this expression as

$$\nabla_p g_\phi \approx \nabla_p \iint_{\phi_a} D(-\mathbf{A}), \tag{3.6}$$

and noting that variations in the gait parameters p affect the gait curve ϕ but not the system's underlying constraint $D(-\mathbf{A})$ (as $D(-\mathbf{A})$ is a property of the system and not the gait), we can convert $\nabla_p g_\phi$ into gradient of an area integral with respect to variations in its boundary. We can then invoke a powerful geometric principle,[2] which states

The gradient of an integral with respect to variations of its boundary is equal to the integral of the [gradient of the boundary with respect to these variations, multiplied by the integrand evaluated along the boundary].

Formally, this multiplication is the *interior product*[3] of the boundary gradient with the integrand,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}), \tag{3.7}$$

which contracts $D(-\mathbf{A})$ (a *differential two-form* [41]) along $\nabla_p \phi$ to produce a differential one-form that can be integrated over ϕ. This formalism will become

[2] The general form of the *Leibniz integral rule* [22].

[3] Not the inner product; the interior product contracts a two-form integrated over areas to a one-form integrated over a path by inserting a vector field as the "first" vector argument of the two-form. See [22] for more details.

important in Chapter 3.2 and Chapter 3.3 when we explore these principles on systems with more than two shape variables; for the two-dimensional shape spaces we consider in this section, the interior product reduces to a simple multiplication between the outward component of $\nabla_p \phi$ and the scalar magnitude of the Lie bracket,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_{p_\perp} \phi)(D(-\mathbf{A})). \tag{3.8}$$

Implementation of the internal pressure: As illustrated in Fig. 3.1, the gradient of the enclosed area with respect to variations in the position of p_i, i.e. $\nabla_{p_i} \phi_a$ in the $e_\parallel$ and $e_\perp$ directions, is the change in triangle's area as p_i moves. Because the triangle's area is always one half base times height (regardless of its pitch or the ratio of its sidelengths), this gradient evaluates to

$$\nabla_{p_i} \phi_a = \begin{bmatrix} e_\parallel & e_\perp \end{bmatrix} \begin{bmatrix} 0 \\ \ell/2 \end{bmatrix}. \tag{3.9}$$

Note that this term matches the right-hand side of (4.37), with only normal motions of the boundary affecting the enclosed area.

3.1.2 Surface Tension from the Distance Metric

The second term in (4.33) takes $\nabla_p s$ as a measure of how variations in the gait affect the cost of executing it, and scales this term by a factor of $\frac{g_\phi}{s}$ to compare how the return on this investment compares to the efficiency of the gait in its

present state. The gradient component of this term can be related to the system's Riemannian metric by first incorporating the arclength calculation from (2.2),

$$\nabla_p s = \nabla_p \oint_\phi (\overbrace{dr^T \mathcal{M}\ dr}^{ds^2})^{\frac{1}{2}} \tag{3.10}$$

and then applying standard calculus operations[4] to arrive at

$$\nabla_p s = \frac{1}{2s} \oint_\phi \Big(2(\nabla_p dr)^T \mathcal{M}\ dr + dr^T (\nabla_p \mathcal{M})\ dr \Big), \tag{3.11}$$

in which the two parts of the integrand respectively measure how changes in the relative positions of the boundary elements and changes in the metric at the underlying points affect the pathlength, and hence the cost of motion.

The $\frac{g_\phi}{s}$ factor (which normalizes the scales of $\nabla_p g_\phi$ and $\nabla_p s$) can be calculated directly from (5.4) and the integral of (2.2). Although the calculation of g_ϕ could, in theory, make use of the area approximation in (5.5), this would be inefficient and impractical: integration of surfaces with arbitrarily complex boundaries requires significantly more computational resources than are needed for line integration around the boundary. Using the true line integral also improves the accuracy of the solution; by continuously recalibrating to the true net displacement, the algorithm avoids compounding any errors introduced by the approximation in (5.5).

Implementation of the surface tension: Each waypoint p_i is at the head of a vector extending from p_{i-1}, such that the dr vector and its gradient in (3.11) can

[4]Namely: Differentiation under the integral sign, chain rule, product rule, and then exploiting the symmetry of $\mathcal{M}$ to consolidate terms.

be taken as

$$dr_i = p_i - p_{i-1} \tag{3.12}$$

$$\nabla_{p_i} dr_i = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \tag{3.13}$$

$$\nabla_{p_{i-1}} dr_i = \begin{bmatrix} -1 \\ -1 \end{bmatrix}. \tag{3.14}$$

For computational simplicity, we evaluate the metric at the center of each segment. This point moves at the mean speed of the segment endpoints, and its gradient with respect to changes in p_i is the mean of its gradient at the endpoints,

$$\nabla_p \mathcal{M}_i = \frac{1}{2} \left(\nabla_{p_{i-1}} \mathcal{M}|_{p_{i-1}} + \nabla_{p_i} \mathcal{M}|_{p_i} \right). \tag{3.15}$$

These gradient relationships are illustrated in Fig. 3.1(b), with the metric represented by its Tissot indicatrix ellipse [32].

3.1.3 Concentration Gradient from Parameterization

Once the gradient descent along (3.2) has given us the image of the optimal gait, a secondary gradient descent along the concentration gradient yields the optimal pacing for the gait. As discussed in §3.1, if we use direct transcription to parametrize the gait, we can simultaneously optimize for the image of the optimal gait, and its

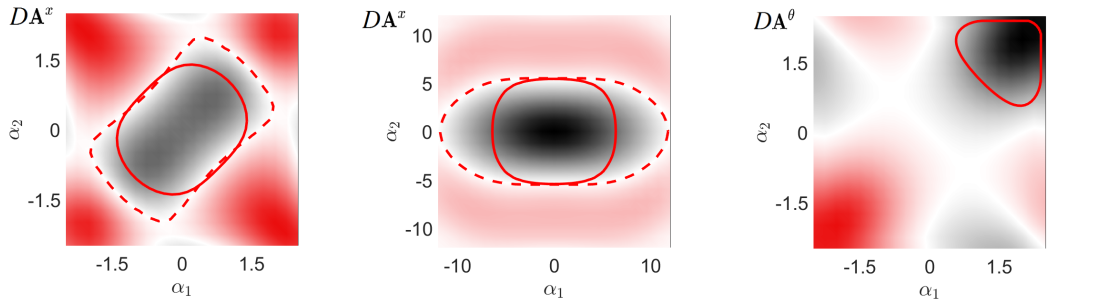

(a) Three-Link Swimmer- Forward Gait (b) Serpenoid Swimmer- Forward Gait (c) Three-Link Swimmer- Turning Gait

Figure 3.2: Optimal gaits found by our algorithm overlaid onto the constraint curvature functions for their corresponding systems. Maximum-displacement gaits, which follow the zero-contours of the corresponding CCF, are indicated by dashed lines. Maximum-efficiency gaits, which are contractions of the zero-contour, are indicated by solid lines. Red regions of the CCF are positive and black regions are negative. The direction of the gait curves is matched to the sign of the region they enclose so that the displacement from each gait is positive.

pacing. In this case, the third term in (4.33), $\nabla_p \sigma$, guides the gait towards pacings (time parameterizations) in which the cycle is executed at a constant rate of power dissipation.

Note that $\sigma = 0$ implies the mapping from τ to $r(\tau)$ gives us a constant power dissipation pacing, which we know is the optimal pacing for executing any gait from [7]. Thus σ serves as an additional cost beyond the best pacing cost of the gait. σ compares the average and local relative rates at which the two length coordinates are changing along the curve. Pushing the gait parameters along the gradient of this stress, $\nabla_p \sigma$, brings the time parameterization into balance with the metric spacing, similarly to how the concentration gradient of soap on the surface of a bubble spreads it into a layer of uniform thickness.

The concentration gradient does not automatically appear in (3.1) or (3.2) be-

cause it is orthogonal to those optimization criteria: tangentially moving points on a curve does not change its length or the region it encloses, and so this gradient is orthogonal to the optimization gradients and acts inside their nullspace (however, see the note below on implementations of this gradient in finite gait parameterizations).

Implementation of the concentration gradient: The local strain energy at a given waypoint corresponds to the square of the difference in the tangential distance from that waypoint to each of its neighbors,

$$\sigma_i = \big((p_{i+1} - p_i)_\parallel - (p_i - p_{i-1})_\parallel\big)^2, \tag{3.16}$$

where $(p_{i+1} - p_i)_\parallel$ refers to the distance between p_{i+1} and $p_i)_\parallel$ along the $e_\parallel$ direction defined at p_i. The gradient of this strain energy with respect to the position of p_i is proportional to the sum of the tangential displacements of the neighboring points relative to p_i,

$$\nabla_{p_i}\sigma \propto \begin{bmatrix} e_\parallel & e_\perp \end{bmatrix} \begin{bmatrix} (p_{i+1})_\parallel + (p_{i-1})_\parallel \\ 0 \end{bmatrix}. \tag{3.17}$$

where $(p_{i+1})_\parallel = (p_{i+1} - p_i)_\parallel$ and $(p_{i-1})_\parallel = (p_{i-1} - p_i)_\parallel$.

Note that in the direct-transcription parameterization, each point under consideration can independently move both tangentially and perpendicularly to the gait curve and that the orthogonality between $\nabla_p\sigma$ and $\nabla_p\frac{g_\phi}{s}$ is thus preserved in the parameterization. Additionally, the spacing provided by $\nabla_p\sigma$ helps ensure the

stability of the optimization by keeping the gait curve from folding over on itself and compromising the gradient calculation.

Other gait parameterizations (e.g., low-dimensional Fourier series) do not necessarily preserve this orthogonality and may lead to the two gradients conflicting with each other. In these cases, the parameterization itself is likely sufficient to prevent the gait curve from becoming degenerate, and the concentration gradient can be left out of the optimization (with the constant-power pacing found by post-processing the output of the curve optimization).

3.1.4 Analysis of the Purcell and serpenoid swimmers

We implemented the optimizer described in the previous section in Matlab, both directly solving the differential equation in (4.33) using `ode45`, and by providing (4.33) as the gradient for the `fmincon` optimizer using the interior-point algorithm. As expected, both implementations converged on the same solutions, with the `fmincon` implementation completing more quickly (on the order of minutes for a modern desktop computer for a gait with 100 parameters), due to its ability to take larger steps through the parameter space. Since the periodicity of the gaits lends itself to a Fourier series parametrization, for the `fmincon` optimizer we generate the direct-transcription waypoints from a lower-order Fourier series parameterization. The low-order restricts the optimizer to simple gaits, preventing neighboring points from crossing each other under the discrete steps taken by fmincon and thus increasing its numerical stability. Another interesting parametrization that could

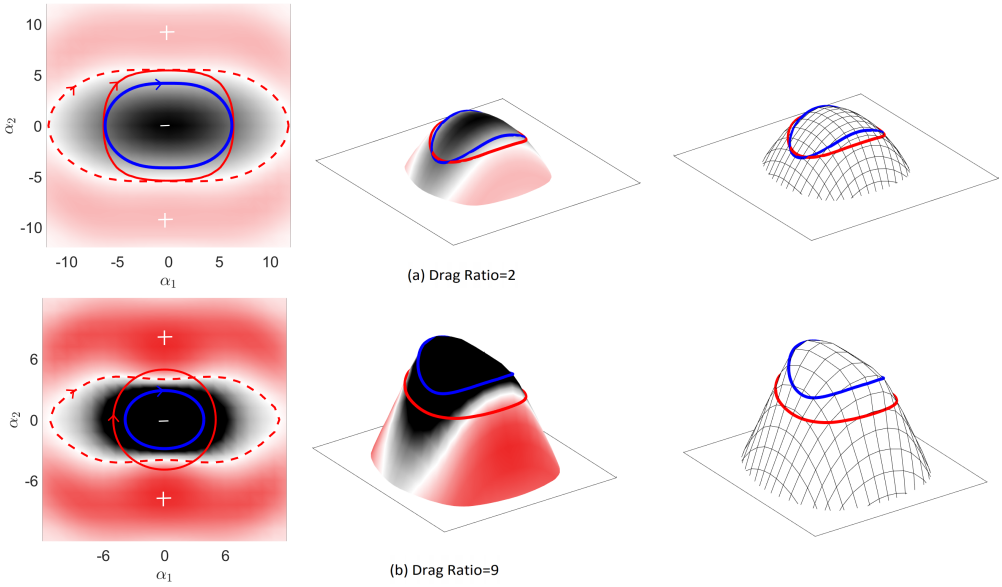

Figure 3.3: Isometric embedding of the serpenoid swimmer's shape manifold in three dimensions for a drag ratio of (a) 2:1 and (b) 9:1. The surfaces in the center and right plots are constructed such that pathlengths over the surface are equal to the metric-weighted pathlengths that capture the cost of executing a gait. The red line indicates the optimal gait when the cost is calculated using the power metric $\mathcal{M}$, and the blue line indicates the optimal gait when the power metric is replaced by an identity metric in the cost calculations (treating the cost as the simple pathlength in the parameter space). Comparing the surfaces in (a) and (b) reveals that the cost manifold for the system with a drag ratio of 9 is more sharply curved than that of the system with a drag ratio of 2. This increased curvature means that the embedding of the manifold is "steeper", and that pathlength thus grows more slowly with increased parameter-radius for the 9-swimmer than the 2-swimmer, leading to a greater difference between the simple pathlength-optimum-stroke and the power-dissipation-optimum stroke for the 9-swimmer.

be incorporated in the future is the one presented in [10], where the motion of each shape variable is parametrized as the sum of a set of compactly supported bump functions added to first-order Fourier series.

We applied the optimizer to the three link system twice, first to find the gait that maximizes the displacement in the x-direction over a single cycle and then to find the maximum-efficiency cycle. The dashed line in Fig. 3.2(a) shows the gait that optimizes the maximum displacement over the cycle for the three-link swimmer. As expected, this gait follows the zero-contour of the height function. The solid line in Fig. 3.2(a) shows the maximum efficiency gait for this system. Because the maximum-efficiency optimizer places a cost on pathlength, this curve gives up the low-yield regions at the ends of the cycle and crosses slightly outside the zero contour. The gaits obtained via this procedure for the three link swimmer match those obtained for maximum efficiency in [28] and, along the x direction, those in [66].

Fig. 3.2(b) illustrates essentially-similar behavior for the serpenoid swimmer: the dashed maximum-displacement gait traces the zero contour, and the solid maximum-efficiency gait captures a more compact area within the sign-definite region. Fig. 3.2(c) shows the gait that produces the most cost-effective rotational motion of the three-link swimmer (the most sign-definite area it could capture on the θ height function while still respecting the constraints on how much the joints can move).

The low-yield region given up at the edges of the zero contour to increase efficiency depends on the anisotropy of the drag experienced by the swimmer in the

lateral and longitudinal directions. Fig. 3.3 shows the shape spaces of serpenoid swimmers with drag ratio of (a) 2:1 and (b) 9:1 as manifolds isometrically embedded in a higher dimensional ambient space such that the cost of executing the gait is the distance traveled on the manifold while executing the gait. Such an embedding lets us accurately visualize and compare the cost of executing different gaits. Similar visualization tools were also used in [17] to study kinematic cartography and in [63] to study two-degrees-of-freedom mechanisms. The greater anisotropy would cause the embedding to have a larger curvature as seen in Fig. 3.3. The reduction in cost associated with the contraction of a gait is smaller on manifolds that are curved more. Thus contractions of zero contour at a drag ratio of two reduces cost (distance traveled on the manifold) more than at a drag ratio of nine, which explains why the optimal gait is closer to the zero contour at a drag ratio of nine than at two.

3.2 Extension to Three Dimensions

For systems with three shape variables, the exterior derivative and local Lie bracket terms from (5.5) each have three components, corresponding to the available pairs

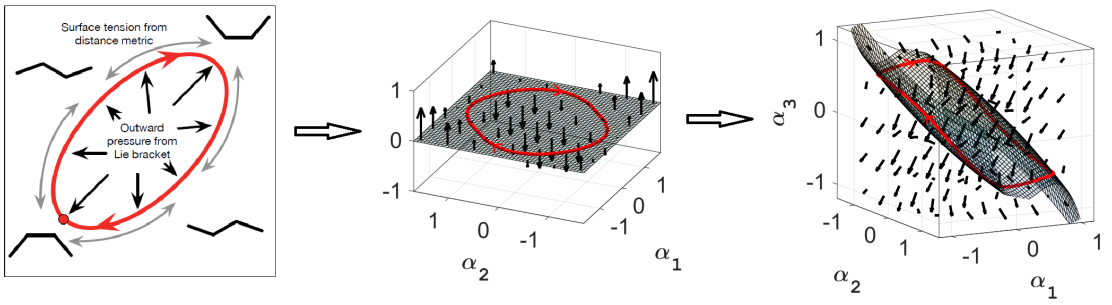

Figure 3.4: Gaits that maximize efficiency enclose the most curvature *of the system dynamics (measured via the curl and Lie bracket of their locomotion dynamics) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). In [59], we showed that for kinematic systems with two shape variables this process is analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble, as shown at left. More generally, the internal pressure can be thought of as being provided by the flux of the curvature passing through a surface bounded by the gait, as in the middle plot. This flux interpretation allows us to boost our geometric framework to three (right plot) or more shape variables.*

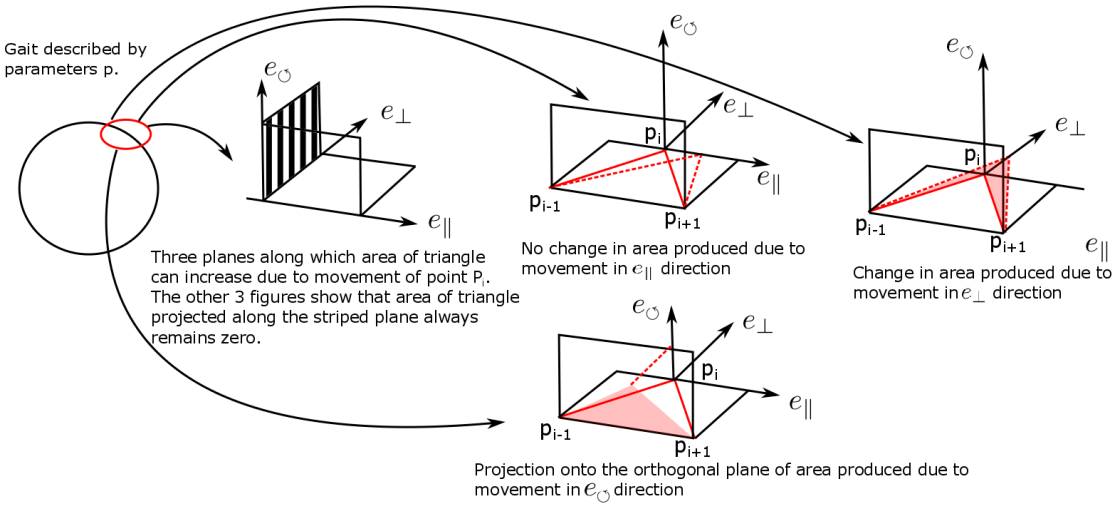

Figure 3.5: Changes in the area enclosed and the direction of change produced by movement along the coordinate directions in the local frame for a three-dimensional shape space. Moving in the tangential direction $e_{\parallel}$ produces no change in the area enclosed, as the area of the triangle given by half the product of base length and height remains the same. Moving in the $e_{\perp}$ direction increases the area along the $e_{\perp} \wedge e_{\parallel}$ plane and moving in the $e_{\circlearrowleft}$ direction increases the area along the $e_{\circlearrowleft} \wedge e_{\parallel}$ plane.

of basis vectors:

$$\mathbf{dA} = \left(\frac{\partial \mathbf{A}_2}{\partial \alpha_1} - \frac{\partial \mathbf{A}_1}{\partial \alpha_2}\right) d\alpha_1 \wedge d\alpha_2 + \left(\frac{\partial \mathbf{A}_3}{\partial \alpha_1} - \frac{\partial \mathbf{A}_1}{\partial \alpha_3}\right) d\alpha_1 \wedge d\alpha_3 + \left(\frac{\partial \mathbf{A}_3}{\partial \alpha_2} - \frac{\partial \mathbf{A}_2}{\partial \alpha_3}\right) d\alpha_2 \wedge d\alpha_3 \quad (3.18)$$

and

$$\sum \big[\mathbf{A}_i, \mathbf{A}_{j>i}\big] = \big[\mathbf{A}_1, \mathbf{A}_2\big] + \big[\mathbf{A}_1, \mathbf{A}_3\big] + \big[\mathbf{A}_2, \mathbf{A}_3\big]. \quad (3.19)$$

The area integral in (5.4) is taken over an oriented surface bounded by the gait.[5]

Employing the same change of coordinates we used for the two-dimensional systems, we can express these two-forms with respect to a local basis in the shape space, in which $e_\parallel$ is tangent to the current gait and $e_\perp$ is normal to the gait in its current plane. We extend this basis into three dimensions with a new vector $e_\circlearrowleft$ that is binormal to the gait trajectory, such that displacing a gait point in this direction would rotate the local surface patch around $e_\parallel$.

As illustrated in Fig. 3.5, extending into the third dimension means that in addition to enclosing extra gait area in the $e_\parallel \wedge e_\perp$ plane by displacing gait points in the $e_\perp$ direction, a gait can also enclose new area in the $e_\parallel \wedge e_\circlearrowleft$ plane by

[5]The existence of a set of such surfaces (the *Seifert surfaces* [67] of the gait) is guaranteed by knot theory; by Stokes' theorem the integral of $\mathbf{dA}$ is the same across all of these surfaces; the $\sum\big[\mathbf{A}_i, \mathbf{A}_{j>i}\big]$ integral may depend on the surface but our minimum-perturbation coordinates [31] make this term small; in any case we are computing the change in this surface with respect to gait parameter variations, and so our algorithm does not need to explicitly identify or integrate over a specific surface.

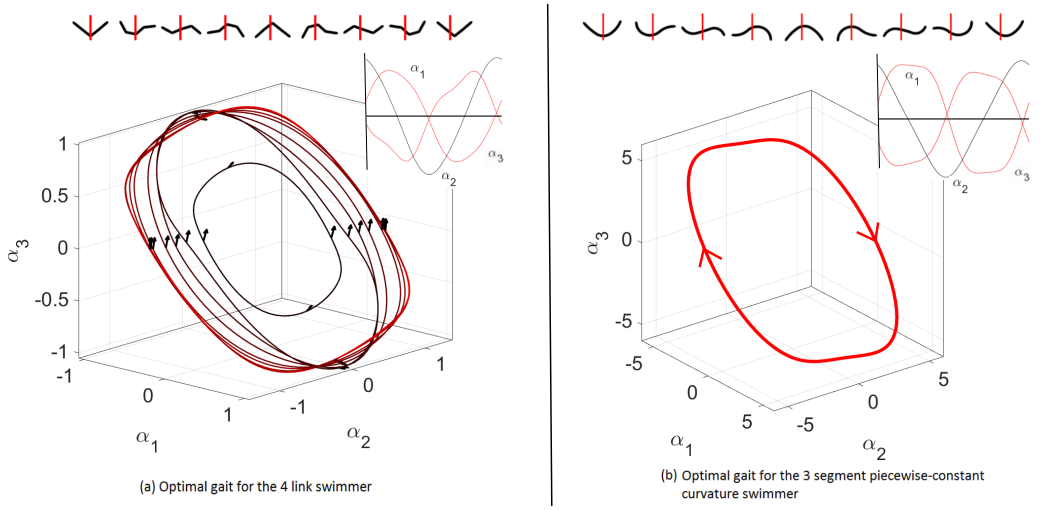

Figure 3.6: Optimal forward gaits for systems with three shape variables. (a) Four-link swimmer. The black ellipse at the center is the gait the optimizer was initialized with, and the outer red rectangle is the optimal gait found by our process. The black arrows show the vector-dual to the constraint curvature, $\overrightarrow{D(-\mathbf{A})}$ at various points. This vector-dual has an approximately constant heading over the region of the shape space explored by the optimizer and the optimal gait thus approximately evolves along a plane. (b) Three-segment piecewise-constant curvature swimmer. Top: The forward progress of the systems as they execute their optimal gait cycles. Both systems move about a tenth of a body length per cycle during these gaits, with the piecewise-constant system requiring less effort for each cycle. Insets: The time-history of the three deformation modes, illustrating a 90° phase shift between the joints on the linked system, and a smaller phase shift (more tightly-grouped peaks) between the segments of the piecewise-constant system.

displacing gait points in the $e_{\circlearrowleft}$ direction. The interior product from (4.36) now preloads the constraint curvature $D(-\mathbf{A})$ with the gradients of gait points in both the normal and binormal directions, such that the gradient of the net displacement with respect to changes in the gait parameters becomes the integral around the gait of the sum of these gradients multiplied by their respective constraint curvature components, such that

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}) \tag{3.20}$$

$$= \oint_{\phi} \Big((\nabla_{p_\perp} \phi) D(-\mathbf{A})_{\parallel\perp} + (\nabla_{p_{\circlearrowleft}} \phi) D(-\mathbf{A})_{\parallel\circlearrowleft} \Big). \tag{3.21}$$

Note that the $\perp\circlearrowleft$ component of $D(-\mathbf{A})$ does not contribute to the gradient of net displacement: no motion of a single point on the gait perimeter can cause its enclosed area to project onto this plane, and so the influence of this term on the gait performance is at most second-order.

Intuitively, this optimization process can be visualized as positioning the gait curve such that it maximizes the flux of a vector field corresponding to $D(-\mathbf{A})$ through a surface bounded by the curve. This vector field $\overrightarrow{D(-\mathbf{A})}$ is formed by associating each plane in the $\mathbb{R}^3$ shape space with its right-hand normal vector (noting that the e_2 basis vector is right-hand normal to the $e_3 \wedge e_1$ plane rather than the $e_1 \wedge e_3$ plane, and adjusting signs accordingly). Motions of the gait points that incorporate more sign-matched area on the $e_{\parallel} \wedge e_{\perp}$ and $e_{\parallel} \wedge e_{\circlearrowleft}$ planes in the two-

form interpretation serve to increase the area of the surface enclosed by the gait and to better-align it with the direction of the flux field, as illustrated in Fig. 3.4(c).[6] For systems with three shape variables, our optimizer is a generalization of [26] from optimal gaits through 2-D slices of high-dimensional spaces to arbitrary paths through these spaces.

3.2.1 Selecting a seed gait

For systems with two shape variables, we know the maximum displacement gait follows the zero contour of the constraint curvature function (CCF). The zero contour of the CCF therefore provides a natural starting point for the optimizer when trying to find the maximum efficiency gait. Analogous to this, in systems with three or more shape variables, we start our optimizer along the zero contour of the CCF projected onto the plane most aligned with the CCF at the extremum of its absolute value.

As discussed above, in systems with three shape variables we can associate a vector field $\overrightarrow{D(-\mathbf{A})}$ with the CCF $D(-\mathbf{A})$. The plane most aligned with $D(-\mathbf{A})$ at the origin is uniquely given by the plane perpendicular to this vector field $\overrightarrow{D(-\mathbf{A})}$

[6] This vector-flux interpretation highlights the point that the *orientation* of the gait in the shape space (determined by the phase offsets between the joint oscillations) is as important as its location in the shape space (the mean and range of the joint motions) when determining the net displacement induced by the cycle.

Simply placing a gait in a region where the signs of all components of $D(-\mathbf{A})$ are the same, as suggested in [64], is neither sufficient nor necessary for generating net motion: A gait in an "all positive" region will produce no net displacement if it is orthogonal to $D(-\mathbf{A})$ (i.e. if $\overrightarrow{D(-\mathbf{A})}$ is in the plane of the gait), and a gait in a "mixed-sign" region can produce displacement proportional to the magnitude of the constraint curvature if it is aligned with $D(-\mathbf{A})$ (i.e., if $\overrightarrow{D(-\mathbf{A})}$ is normal to the surface).

at the origin. The zero contour of $D(-\mathbf{A})$ projected onto this plane forms the seed gait for our optimizer in systems with three shape variables.

3.2.2 Analysis of the four-link swimmer

We used the optimizer to find the most efficient forward gait for a four-link (three-joint) swimmer. The optimal gait is shown in Fig. 3.6(a). The optimal gait for the four-link swimmer is 50% more efficient than a three-link swimmer with same total length and same drag coefficients. (Three link swimmer efficiency is 0.11, and that of the four-link swimmer is 0.16), but it is still less than the efficiency of a serpenoid swimmer with two mode shapes (0.24).

From Fig. 3.7, we can see that the optimal gait for the four-link swimmer lies very close to the $\alpha_3 = -\alpha_1$ plane, such that these joints oscillate almost 180° out of phase with each other and cross each other at zero angle. The α_2 oscillation is at approximately 90° phase difference to the other two joint motions and is slightly higher amplitude, tilting the gait within this plane.

3.2.3 Analysis of three segment piecewise continuous curvature swimmer

We used the optimizer to find the most efficient forward gait for a system with three piecewise-continuous curvature segments. The optimal gait is shown in Fig. 3.6(b). This gait is similar to the one we found for the linked system, but it has a smaller

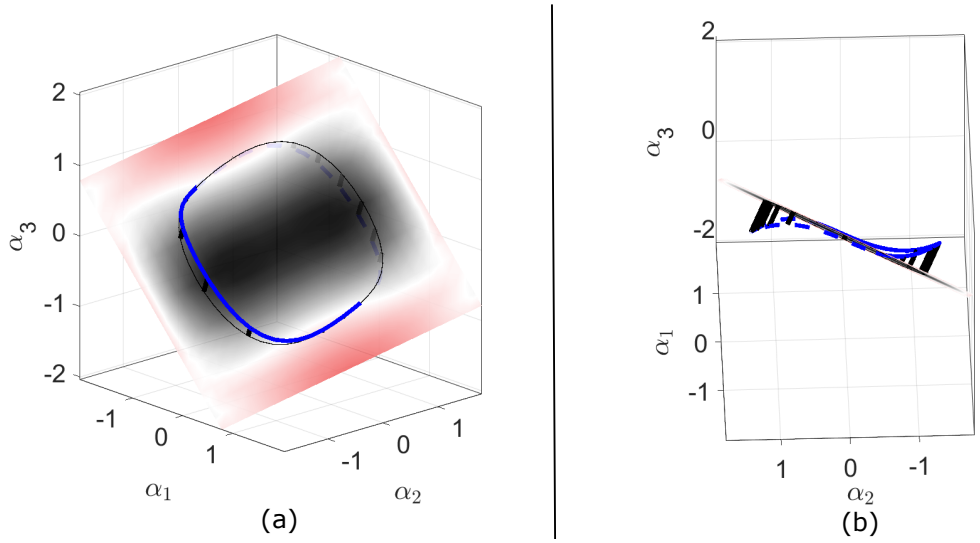

Figure 3.7: Projection of the constraint curvature function for the 4-link swimmer onto the plane most aligned with it at its center. The blue line indicates the optimal forward gait for the 4-link swimmer. (a) The projection of the optimal gait onto this plane is a contraction of the zero contour in that plane. (b) The optimal gait for the five-link swimmer only deviates slightly from this plane, hence the zero contour of the constraint curvature function projected onto this plane provides a good seed gait for the optimizer.

phase shift between adjacent actuators, and so is rotated around the α_3 axis by approximately 45°. The efficiency of this gait is 0.18, which is better than the four-link swimmer but still falls short of the efficiency of the serpenoid swimmer with two shape modes.

3.3 Extension to n dimensions

The gradient of gait displacement for systems with n-dimensional shape spaces has a a similar form to the three-dimensional formulation in (3.21). The key difference is that there are now $n-2$ "binormal" directions[7] in which gait points can be displaced, making the displacement gradient

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} \left((\nabla_{p_\perp} \phi) D(-\mathbf{A})_{\parallel\perp} + \sum_{i}^{n-2} (\nabla_{p_{\perp_i}} \phi) D(-\mathbf{A})_{\parallel\perp_i} \right). \quad (3.22)$$

Note that because the interior-product formulation excludes the influence any components of $D(-\mathbf{A})$ that do not include $e_\parallel$, the number of components in (3.22) goes up linearly with n, even though the number of independent planes (and thus components of $D(-\mathbf{A})$) scales quadratically as $\frac{n(n-1)}{2}$, as illustrated in Fig. 3.8.

The vector-flux analogy that we made in three dimensions becomes a "patch-flux" analogy as we move into higher dimensions: The basis areas for two-forms no

[7] Outside of three dimensions, "rotation around a line segment" is not a well-defined operation, so we no longer call these directions "rotational."

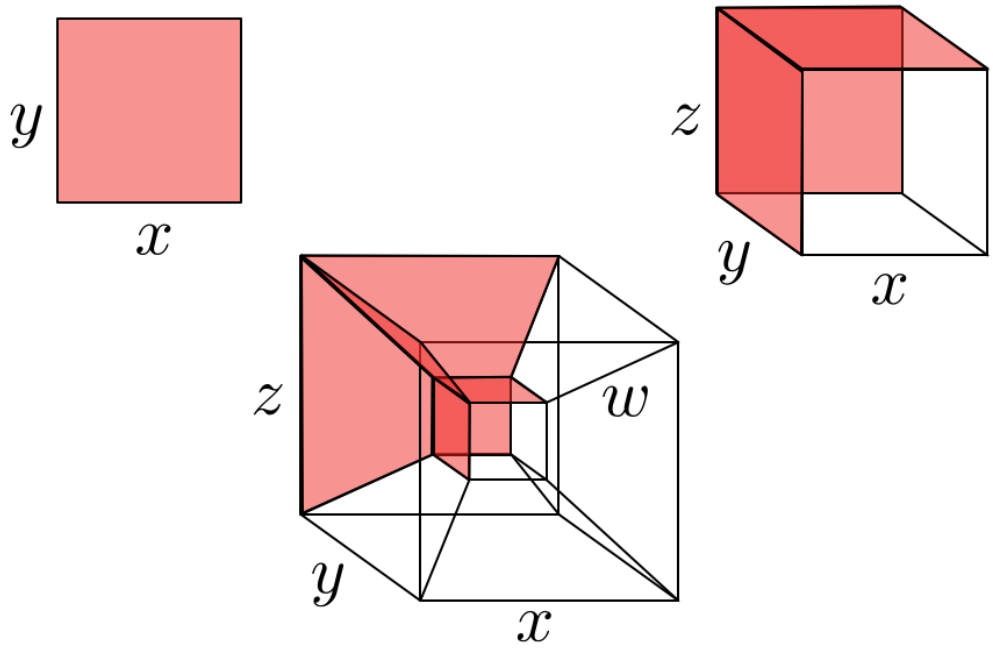

Figure 3.8: The number of components of the constraint curvature in n-dimensional space is equal to the number of non-parallel two-dimensional faces on an n-cube: one for the 2-cube, or square; three for the standard 3-cube; and six for the 4-cube (or "tesseract"). (Faces on the tesseract are considered parallel if they are formed from the same pair of basis vectors, e.g., all xw faces are parallel to each other.) For higher-dimensional spaces, the number of independent two-dimensional faces continues to scale as $\frac{n(n-1)}{2}$.

longer have unique normal vectors, so we cannot directly map $D(-\mathbf{A})$ to a vector field. We can, however, still take the area integral of $D(-\mathbf{A})$ as an area integral over a surface, and each change of the gait boundary adds an infinitesimal patch to the edge of this surface, each of which contributes to the net integral based on its alignment with the basis areas of the space and the values of $D(-\mathbf{A})$ in those basis-area directions.

3.3.1 Selecting a seed gait

In systems with more than three shape variables, the problem of finding a plane most aligned with the $D(-\mathbf{A})$ can be more generally restated as finding two unit vectors $v, w \in T_{\mathbf{0}}R$ that maximize the value of $D(-\mathbf{A})|_{r=\mathbf{0}}(v, w)$. We shall refer to the constraint curvature function at the extremum of its absolute value, $D(-\mathbf{A})|_{r=\mathbf{0}}$ as $D(-\mathbf{A})_{\mathbf{0}}$. We can associate a real skew-symmetric matrix B to $D(-\mathbf{A})_{\mathbf{0}}$ such that

$$D(-\mathbf{A})_{\mathbf{0}}(v, w) = v^T B w \quad \forall v, w \in T_{\mathbf{0}}R \tag{3.23}$$

We know that the eigenvalues of real skew symmetric matrices are purely imaginary. Let $i\lambda$ be the largest eigenvalue and z be the unit eigenvector associated with it. let $z = x + iy$ where x and y are real column vectors. Therefore,

$$Bz = i\lambda z \tag{3.24}$$

$$B(x + iy) = -\lambda y + i\lambda x \tag{3.25}$$

Therefore, $x^T B y = \lambda \|x\|_2^2$ and $y^T B x = -\lambda \|y\|_2^2$. Since B is skew-symmetric, $x^T B y = -y^T B x$, therefore $\|x\|_2 = \|y\|_2 = \frac{1}{\sqrt{2}}$ as z is a unit eigenvector.[8] Therefore if $v = \sqrt{2}x$ and $w = \sqrt{2}y$, then $D(-\mathbf{A})_{\mathbf{0}}(v, w) = \lambda$. Since λ is the magnitude of the largest eigenvalue of B, this is the largest value $D(-\mathbf{A})_{\mathbf{0}}$ can attain by acting on two unit vectors.

[8]Note that outside of the case when $D(-\mathbf{A}) = 0$ on the whole region, x cannot be equal to y as that would imply $x^T B x = -x^T B x$, hence $x = y = 0$ which would mean the largest eigenvalue of B is zero which would mean $D(-\mathbf{A})$ is zero at the at the extrema of its absolute value.

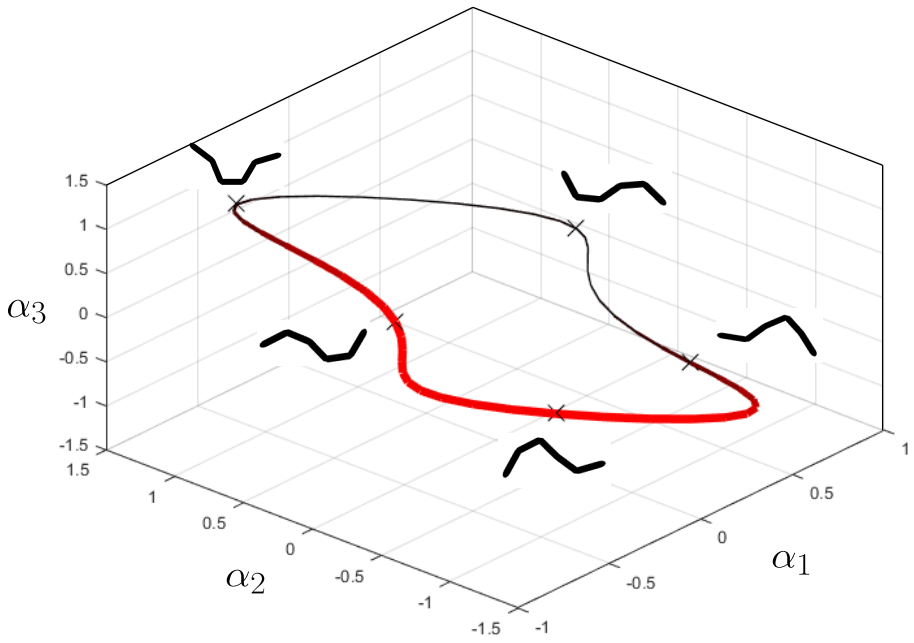

Figure 3.9: Optimal forward gait for the five-link swimmer. The three axes represent the first three joint angle. The value of the fourth joint angle is given by the color and thickness of the line. Thin black lines indicates sections of the gait with low values of α_4 and thick red lines indicate sections of the gait with high α_4.

Thus the plane most aligned with the constraint curvature function at the origin is the plane spanned by vectors v and w such that $z = v + iw$ is an eigenvector corresponding to the largest eigenvalue of the skew symmetric matrix B associated with $D(-\mathbf{A})$ at the origin. We project $D(-\mathbf{A})$ onto this plane and use the zero contour thus obtained as the seed gait for our optimizer.

3.3.2 Analysis of the five-link swimmer

We used the optimizer to find the most efficient forward gait for a five-link (four-joint) swimmer. The optimal gait is shown in Fig. 3.9. The three axes represent the first three joint angle. The value of the fourth joint angle is given by the color and thickness of the line. Thin black lines indicates sections of the gait with low values of α_4 and thick red lines indicate sections of the gait with high α_4. The optimal gait for the five-link swimmer has an efficiency of 0.21

3.3.3 Comparison of different swimmers

In Fig. 3.10, we present a comparison of the efficiency of optimal gaits found for various systems of up to four shape dimensions. Fig. 3.10 shows a comparison of the cost and displacement produced by executing these gaits. For two shape modes, the serpenoid system outperforms the jointed and piecewise-constant systems, but its efficiency improves only slightly if we add a second pair of (double-spatial frequency) shape modes. The jointed and piecewise-constant systems both exhibit a pattern where moving from two to three shape degrees of freedom allows them to lower their cost of motion while increasing the net displacement, and then moving from three to four shape modes significantly increases the net displacement while increasing the cost of the motion only slightly. at any level of articulation, the piecewise-constant system outperforms the jointed system and at four degrees of freedom approaches the serpenoid efficiency.

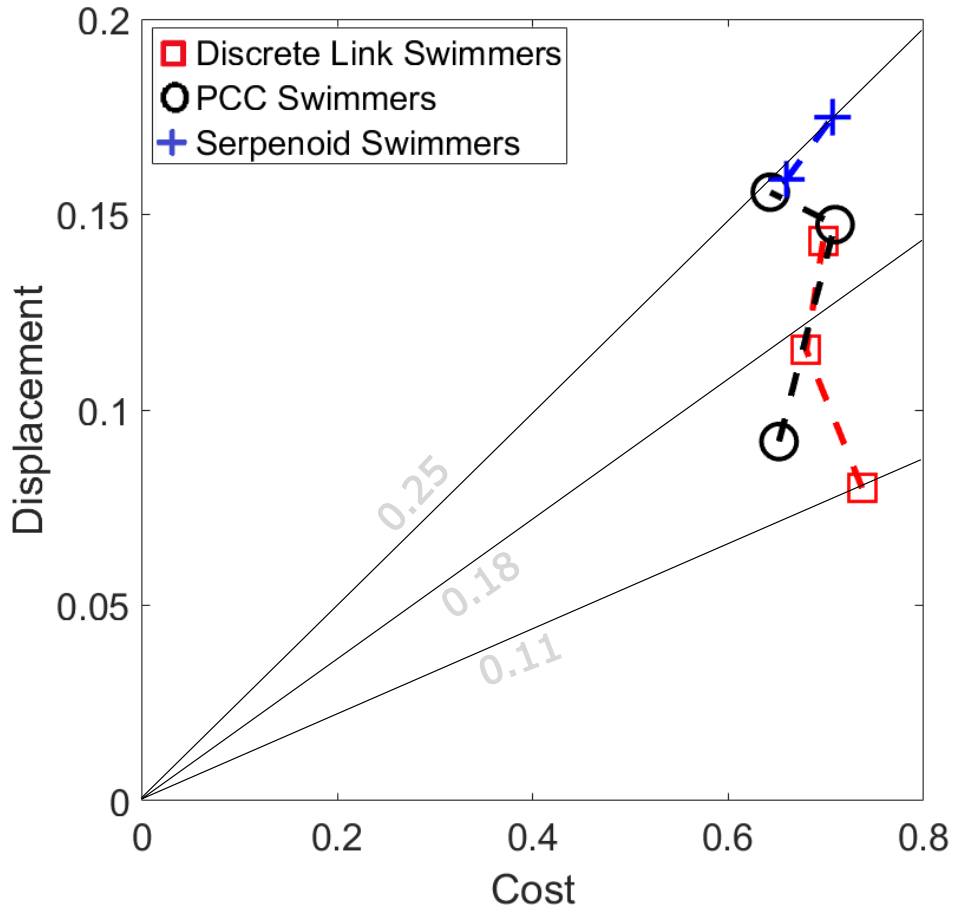

Figure 3.10: Comparison of the efficiencies of different swimmers. Red squares represent discrete link swimmers, black circles represent piecewise continuous swimmers, and blue plus signs represent serpenoid swimmers. The number next to each symbol indicates the shape variables of the system. Displacement produced by the most efficient gait for each swimmer is given by the ordinate value, and the cost of executing the gait is given by the abscissa value of the symbol representing the swimmer, and their efficiencies are the slopes of the lines connecting them to the origin. The serpenoid swimmers are the most efficient, followed by the piecewise swimmers and then the discrete link swimmers. A movie presenting a side-by-side comparison of these gaits is included in the Supplementary Material.

3.4 Simultaneous Design and Gait Optimization

A common approach to finding the optimal value of a design variable for a locomoting system is via a nested optimization in which the outer loop optimizes the design variable, and the inner loop finds the optimal gait for each geometry considered in the outer loop. Our variational framework can be used to unify the design and gait optimization process for drag dominated systems.

By introducing the design variables as extra shape variables with the constraint that these variables are constant within the gait cycle (i.e., that the gaits trace out level sets in the design parameters), we can use our optimizer to simultaneously find the optimal value for the design variables and the optimal gait for that design. For example, if we wanted to design a three-link swimmer of a specified total length, the optimal ratio between the length of the middle link to the length of the outer links would be the one at which the optimal gait at that ratio is more efficient than the optimal gaits for other ratios.

To include design variables in our optimization, we treat them as pseudo-shape variables. Suppose $\beta = (\beta_1 \cdots \beta_m)$ are our design variables. We redefine our shape variables to be $r = (\alpha_1, \cdots, \alpha_n, \beta_1, \cdots, \beta_m)$. We then restrict our optimizer to only move points in ways that would keep the value of the design variables constant around the gait by imposing the constraint $\beta_i|_{p_j} = \beta_i|_{p_k}$, for each design parameter β_i, across all points p_j and p_k.

By simultaneously optimizing the design and control variables, our procedure avoids having to compute the optimal gait at each intermediate set of design val-

ues, and thus requires many fewer iterations than a nested optimization scheme. If this simultaneous optimization were incorporated into a numerical optimizer without the benefit of our geometric construction of the gradient, then the design variables would compound the dimensionality of the control space, and the ensuing complexity would not be mitigated by the locality of the geometric process.

Identifying an optimal link-length ratio for the Purcell swimmer

We used our optimizer to find that the optimal ratio between the length of the middle link to the length of the outer links for the Purcell swimmer is 0.80, which is close to the value of 0.75 obtained by nested optimization in [66]. The efficiency of the optimal gait we found at the optimal link ratio is 0.12, which matches the efficiency found for the optimal gait on the 0.75-ratio system in [66] (i.e., our results are equivalent to previous results up to the precision of our gait-displacement ODE solver).

The steps[9] that our optimizer takes from a (deliberately non-optimal) seed gait and equal link lengths to the optimal gait at an optimal ratio are illustrated in Fig. 3.11. Note that because we are simultaneously stepping along the gait and design gradients, the system does not monotonically approach the optimal link ratio, but instead moves to design variables that make the best use of the current gait geometry; as the gait geometry moves from the sub-optimal ellipse to the more efficient round shape, the design variable settles to its optimal value.

[9]Every third iteration of `fmincon`, using our provided gradient.

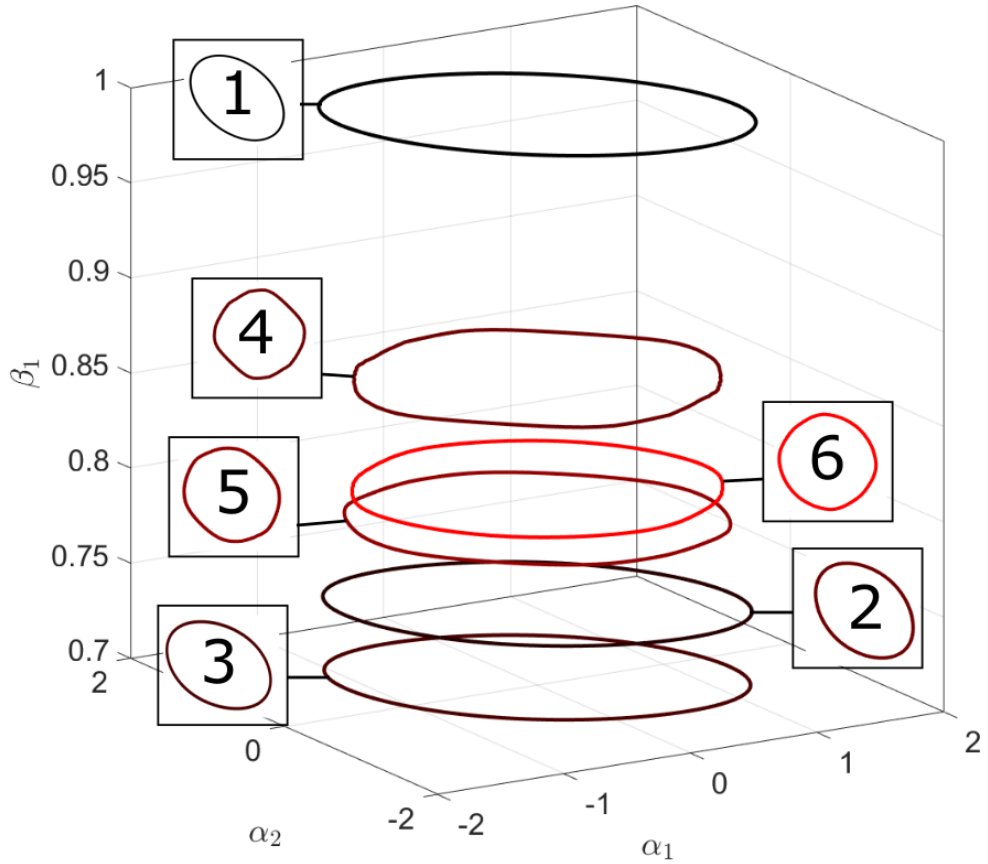

Figure 3.11: This figure shows the process of finding the optimal link length ratio and the maximum efficiency gait for the Purcell swimmer simultaneously. We seed our optimizer at the black gait (1), and the red gait (6) is the optimal gait. The inset figures provide a top view of the gaits at each iteration, and the numbers show the steps through which the gait and design variables evolve

3.5 Conclusion

In this chapter, we consolidate upon and extend prior geometric insights about locomotion into a set of geometric principles that govern the shape of optimally-efficient gaits for drag dominated systems. We formally encode these principles in a set of geometric expressions that together make up the gradient of the gait efficiency with respect to variations in the gait trajectory. We use this gradient in a gradient-descent solver to find optimal gaits, but more fundamentally, *the gradient geometrically describes the dynamics underlying any other gait optimization algorithm applied to the system.*

For systems with two shape variables, the dynamics of this solver are analogous to those of a soap bubble, with the Lie bracket providing an "inflating pressure" to the trajectory and the Riemannian metric on the shape space contributing "surface tension" that halts growth of the cycle in the face of diminishing returns, and a "concentration gradient" that provides a power-optimal pacing along the gait. Together, these elements drive the gait cycle to a "comfortable stride" that converts shape change effort into net displacement with optimal efficiency.

By extending the gradient calculation to systems with three shape variables we see that dynamics of our solver generalize to maximizing flux through an oriented surface. For systems with three shape variables, the constraint curvature functions can be visualized as vector flux because each surface element has a unique normal.

In systems with more than three shape variables, each surface element no longer has a unique normal direction associated with it, and we thus drop the "vector

flux" analogy for these higher-dimensional systems, but can preserve the idea of two-form flux passing through the surface elements that make up the "interior" of the gait.

We demonstrated this variational principle in operation on a number of test systems in viscous-dominated environments, including Purcell's three-link swimmer (a standard minimal template for locomotion modeling) and a serpenoid swimmer (a model widely used in studies of animals and snake robots). In the lower-dimensional cases, the optimal gaits found by our approach match those previously found by exhaustive optimizations of the gait cycles [66], and in the higher-dimensional cases, the optimizer allowed us to efficiently explore a space of candidate swimming morphologies.

We also presented how the framework can be used to simultaneously optimize design and control variables for locomoting systems. We demonstrated this by finding the optimal link length ratio for the Purcell swimmer.

In the context of related works, our framework can be viewed as a "macroscopic" extension of the Lie-bracket control schemes in [50, 52]. This extension is significant because it uses the geometry of the systems' dynamics to identify the *amplitudes* of their most efficient gait cycles (in addition to the useful sets of phase-couplings identified in previous work). For our example systems swimming at low Reynolds number, the two-dimensional "soap-bubble" analysis can be seen as an instantiation of the boundary-value problem suggested in [3], for which we have analytically constructed a gradient from the curvature of the constraints, and the higher-dimensional portions of a work to be a generalization of this princi-

ple. These extensions and generalizations of previous work in the literature have both been enabled by our calculation of a minimum-perturbation body frame [34], which significantly increases the accuracy of Lie-bracket approximations to large motions.

Chapter 4: Swimmers with Elastic Tails

In this chapter, we extend the framework presented in chapter 3 to identify optimal gaits for swimmers with a passive elastic joint. Drag-dominated swimmers with passive elements were studied in [15, 38, 42, 49, 56]. Of these works, [42, 56] and [49] are most relevant to this chapter, as they discuss the motion of swimmers with a harmonically-driven active joint and a passive joint. The analysis in [56] is particularly relevant, where using perturbation expansion, explicit expressions for leading order solutions were derived for harmonic input oscillations, and the optimal swimmer geometry was obtained for the Purcell swimmer. Unlike in [56], we do not restrict our input to simple harmonic oscillations, and use a higher order representation of the system dynamics.

The ways in which the dynamics of swimmers with passive elastic joints differ from the dynamics of fully actuated swimmers are:

1. Because of the coupling between the actuated and passive joints, the passive swimmers can execute only a subset of the gait kinematics the fully actuated swimmers can execute.
2. This coupling between the actuated and passive joints also endows each gait achievable by the passive swimmer with a unique pacing. Hence, the pacing cost cannot be minimized separately from the kinematic cost.

As illustrated in Fig. 1.2, we address the two problems introduced by the presence of passive elastic joints: we first use frequency domain analysis to analytically approximate the motion of the unactuated joint in response to the motion of the actuated joint. We then combine this frequency-space analysis with elements of the geometric framework introduced in [61] to construct a gradient-descent algorithm that identifies optimal gaits and the pacing associated with these gaits for passive swimmers. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

4.1 Frequency Domain Analysis

The key difference in the dynamics of a swimmer with an elastic joint when compared to a fully actuated swimmer is the coupling of the motion of the actuated and unactuated joint. In this section, we explore this difference further and present a way of accurately approximating the motion of the unactuated joint from the motion of the actuated joint using frequency domain analysis. The method of linearizing the passive dynamics to obtain approximate limit cycles presented in this section is in the same vein as the limit cycle analysis presented in [15], where a two-link system with static separation between centers of mass and buoyancy was studied.

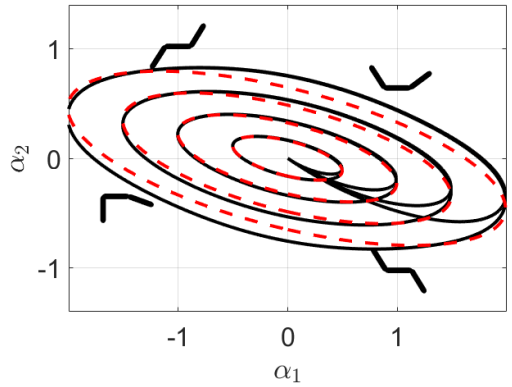

(a) Purcell swimmer

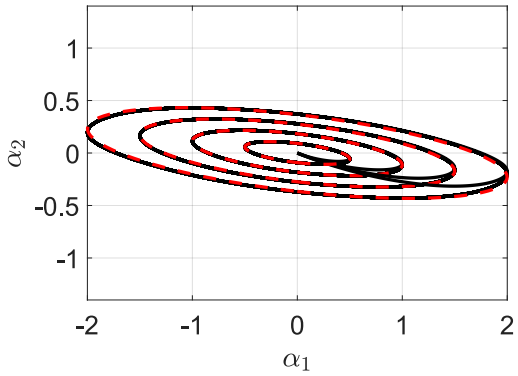

(b) T-link swimmer

Figure 4.1: Comparison of the exact stroke limit cycles with the shape predicted by frequency domain analysis. Each black solid line represents the motion of the swimmer when the input to the actuated joint is a sinusoidal wave. We can see the motion converges to a limit cycle. The red dashed line represent the shape of the limit cycles predicted by frequency domain analysis presented in §4.1. The cartoons in the background show how the Purcell swimmer looks like at different points of the shape space.

4.1.1 Dynamics of the passive elastic joint

As discussed in [33, 34], the mapping between joint velocities and torques on the joint in the fully actuated swimmers is encoded in the metric calculated in (2.11),

$$\tau = \mathcal{M}(r)\dot{r}. \tag{4.1}$$

Since the systems considered in this paper have only one active and one passive joint,

$$\begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} = \mathcal{M}(\alpha_1, \alpha_2) \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \tag{4.2}$$

In the case of the swimmers with an elastic joint, because the actuation in the second joint is replaced by an elastic element with stiffness k, the torque τ_2 is always equal to $-kr_2$, i.e.,

$$\begin{bmatrix} \tau_1 \\ -k\alpha_2 \end{bmatrix} = \mathcal{M}(\alpha_1, \alpha_2) \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \tag{4.3}$$

The first equation in this system of equations,

$$\tau_1 = \mathcal{M}_{11}\dot{\alpha}_1 + \mathcal{M}_{12}\dot{\alpha}_2, \tag{4.4}$$

thus relates the torque in the actuated joint to the motion of the joints and can be used to calculate the torque required to effect any feasible motion. The second

equation in this system of equations,

$$-k\alpha_2 - \mathcal{M}_{22}\dot{\alpha}_2 = \mathcal{M}_{21}\dot{\alpha}_1 \tag{4.5}$$

or equivalently,

$$-\frac{k}{\mathcal{M}_{21}}\alpha_2 - \frac{\mathcal{M}_{22}}{\mathcal{M}_{21}}\dot{\alpha}_2 = \dot{\alpha}_1 \tag{4.6}$$

encodes the dynamics of the passive elastic joint in terms of the active joint, and thus defines the space of feasible motions.

Since the joint α_1 is assumed to be the actuated joint, we have full control of $\dot{\alpha}_1$. The value of $\mathcal{M}$ depends on α_1 and α_2. Its dependence on α_1 and α_2 conveys how the shape of the robot affects the effort required to move the joints.

If we consider gaits that are relatively small oscillations of shape, we can approximate the value of $\mathcal{M}$ to be constant throughout the gait. This assumption necessarily introduces errors in our prediction of the motion of the passive joint when the amplitude of input to the active joint is large. For both the Purcell and T-link swimmers with passive elastic joints, assuming the value of $\mathcal{M}$ to be constant does not introduce significant errors for gaits of amplitude up to 1.5 radians.

In Fig. 4.1, we illustrate the distortions caused in the shape of the limit cycles when we assume $\mathcal{M}$ to be constant throughout the shape space. Each solid black line represents the motion of the full swimmer model when the input to the actuated joint is a sinusoidal wave, and the system starts with both angles at zero. There is a transient term that dominates before the system reaches the limit cycle.

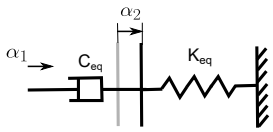

Figure 4.2: Spring damper system whose dynamics are equivalent to the passive dynamics of the Purcell swimmer with a passive elastic joint.

The red dashed line represents the shape of the limit cycles predicted when we assume $\mathcal{M}$ is constant (we describe the calculation of the limit cycle in the next subsection).

4.1.2 Transfer function analysis

In this subsection, our goal is to obtain an analytical approximation of the response of the passive joint to input oscillations of the actively controlled joint. We assume $\mathcal{M}$ to be constant throughout the gait, which makes (4.6)— the equation that describes the dynamics of the passive elastic joint— a linear first order differential equation and thus well suited to frequency domain analysis.

We use the Laplace transform on (4.6) to obtain the dynamics of the passive elastic joint in the frequency domain as

$$K_{eq}\mathcal{L}(\alpha_2(t)) + C_{eq}s\mathcal{L}(\alpha_2(t)) = s\mathcal{L}(\alpha_1(t)) \tag{4.7}$$

where $K_{eq} = \dfrac{-k}{\mathcal{M}_{21}}$ and $C_{eq} = \dfrac{-\mathcal{M}_{22}}{\mathcal{M}_{21}}$. These substitutions reveal that the

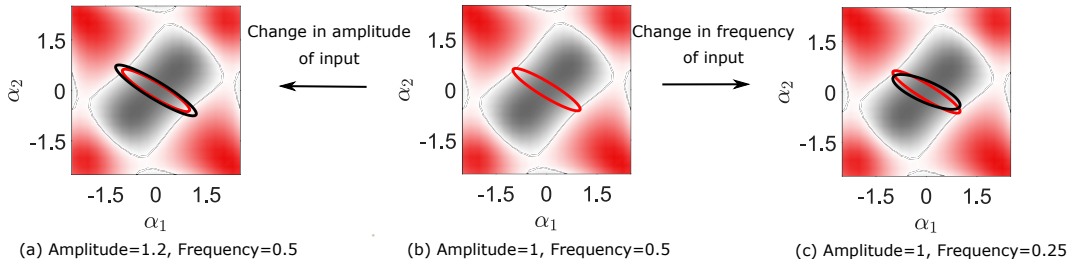

Figure 4.3: Changes in the shape of a gait resulting from changes to amplitude and frequency of input oscillation. (a) The gait resulting from a change in amplitude of the nominal input (black) is a scaled version of the gait corresponding to the nominal input actuation (red). (b) The gait resulting from nominal input actuation. (c) The gait resulting from a change in the frequency of the nominal input (black). A change in the frequency of the nominal input leads to a change in both the amplitude and phase of the response of the passive joint.

dynamics of the passive joint resemble those of a massless particle attached to a fixed base through a spring and being driven through a damper by a position trajectory α_1 as shown in Fig. 4.2. Note that the damping coefficient is completely dependent on the physics of the interaction between the swimmer and the fluid.

We can rewrite (4.7) as a transfer function relation between the active joint α_1 and passive elastic joint α_2

$$\mathcal{L}(\alpha_2(t)) = H(s)\mathcal{L}(\alpha_1(t)) \tag{4.8}$$

where $H(s) = \dfrac{s}{(C_{eq}s + K_{eq})}$ is the transfer function that encodes the response of the passive elastic joint to osciallations of the active joint.

Equation (4.8) tells us how inputs to the controlled joint, α_1 are mapped to the response of the passive joint in the frequency domain. In order to find the

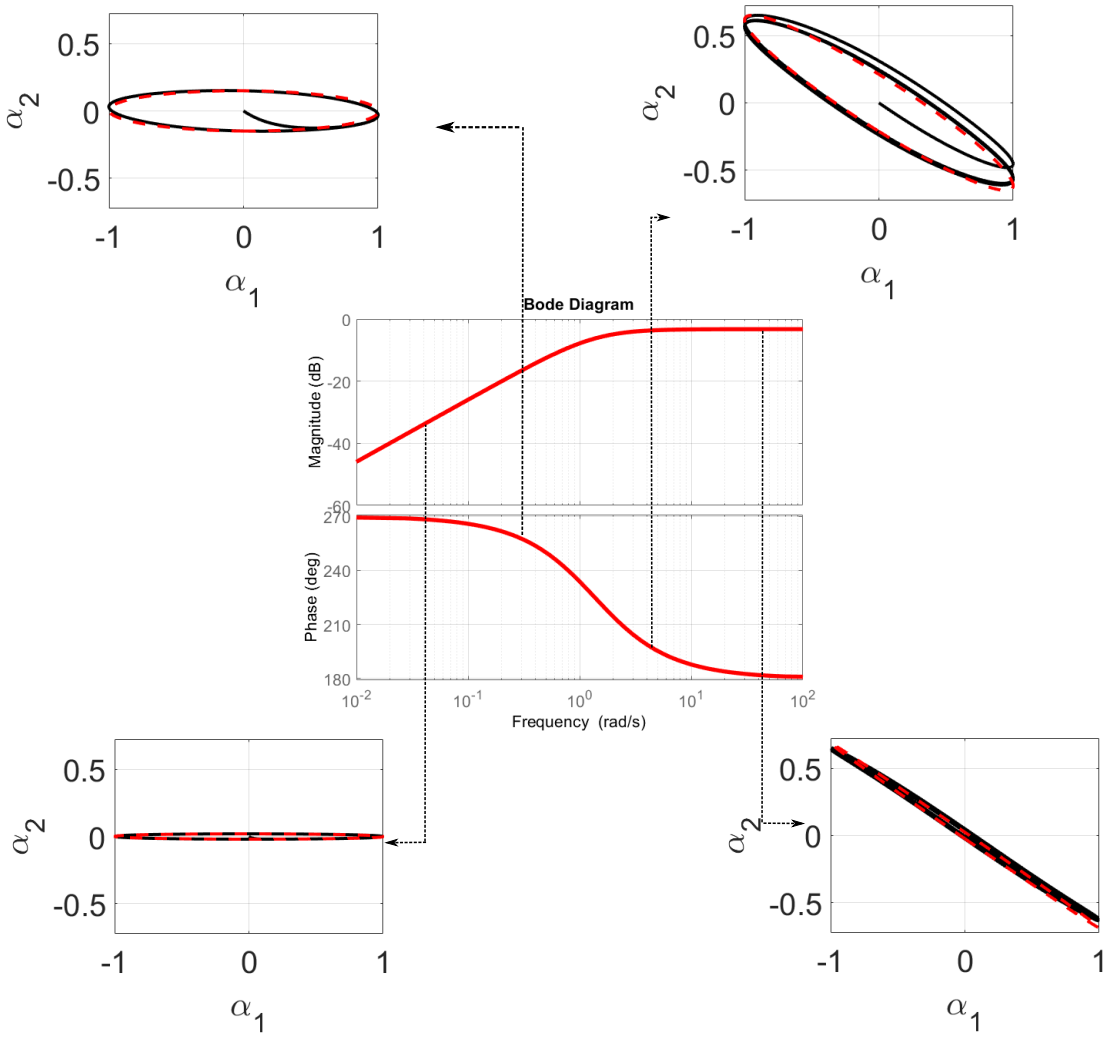

Figure 4.4: Bode plot of the transfer function relating the output of the passive joint to the input of the actuated joint in the Purcell swimmer with a passive elastic tail. The inset figures show how the periodic orbits corresponding to gain and phase at certain frequencies look like in the shape space. We can see that actuation at very low frequencies leads to gaits that enclose very little surface area due to the amplitude of the passive joint being low, while at very high frequencies the surface area enclosed by the gait is low due to the almost 180 degree phase shift between the oscillations of the actuated and passive joints.

response of the passive joint to a sinusoidal oscillation of the actuated joint, we let $\alpha_1(t) = \sin(wt)$, then using (4.8), we obtain

$$\mathcal{L}(\alpha_2(t)) = \overbrace{\frac{}{(C_{eq}s + K_{eq})}}^{H(s)} \overbrace{\frac{\omega}{(s^2+\omega^2)}}^{\mathcal{L}(\alpha_1(t))} \tag{4.9}$$

$$= \underbrace{\frac{A_1}{(C_{eq}s + K_{eq})}}_{\text{Transient term}} + \underbrace{\frac{A_2 s + A_3}{(s^2+\omega^2)}}_{\text{Phase shifted sine wave}} \tag{4.10}$$

where A_1, A_2 and A_3 can be obtained by equating the two expressions for $\mathcal{L}(\alpha_2(t))$ in (4.9) and (4.10) as

$$s\omega = A_1(s^2+\omega^2) + (A_2 s + A_3)(C_{eq}s + K_{eq}), \tag{4.11}$$

and equating the coefficients of powers of s on each side to extract a system of three equations,

$$\begin{bmatrix} 1 & C_{eq} & 0 \\ 0 & K_{eq} & C_{eq} \\ \omega^2 & 0 & K_{eq} \end{bmatrix} \begin{bmatrix} A_1 \\ A_2 \\ A_3 \end{bmatrix} = \begin{bmatrix} 0 \\ \omega \\ 0 \end{bmatrix}, \tag{4.12}$$

which can be easily solved to obtain A_1, A_2 and A_3 as

$$\begin{bmatrix} A_1 \\ A_2 \\ A_3 \end{bmatrix} = \begin{bmatrix} 1 & C_{eq} & 0 \\ 0 & K_{eq} & C_{eq} \\ \omega^2 & 0 & K_{eq} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \omega \\ 0 \end{bmatrix}. \tag{4.13}$$

If our actuated joint follows a trajectory given by $\alpha_1(t) = B_1 \sin(\omega t)$, after reaching the steady state periodic orbit, the unactuated joint, therefore, follows the trajectory

$$\alpha_2(t) = B_1 A_2 \cos(\omega t) + B_1 A_3 \sin(\omega t). \tag{4.14}$$

If our actuated joint follows a trajectory given by $\alpha_1(t) = B_1 \sin(\omega t) + B_2 \cos(\omega t)$, after reaching the steady state periodic orbit, the unactuated joint follows the trajectory

$$\begin{aligned} \alpha_2(t) &= B_1(A_2 \cos(\omega t) + A_3 \sin(\omega t)) \\ &\quad + B_2(-A_2 \sin(\omega t) + A_3 \cos(\omega t)). \end{aligned} \tag{4.15}$$

Note that the value of A_2 and A_3 depend on the value of ω, the frequency of the sinusoidal input to the actuated joint. The image of the gait in the shape space is thus coupled to the pacing of the input to the actuated joint.

This coupling between the image of the gait in the shape space and the input to the actuated joint is illustrated in Fig. 4.3. Fig. 4.3(b) shows the shape of the gait resulting from a sinusoidal oscillation of the actuated joint of amplitude 1 and frequency 0.5. Fig. 4.3(c) shows the effect of decreasing the frequency of oscillation on the shape of the gait, with a slower oscillation leading to a weaker response (in terms of magnitude) from the passive joint. Fig. 4.3(a) shows the effect of increasing the amplitude of oscillation on the shape of the gait. As we have linearized the dynamics of the passive joint, a change in amplitude without a change in frequency produces a scaled version of the original gait.

The Bode plot of the transfer function H in (4.8), presented in Fig. 4.4, illustrates the specific nature of the relationship. We can see that if the frequency of input to the controlled joint is very high, the response of the passive elastic joint is phase shifted relative to the actuated joint by almost 180 degrees, resulting in a gait with a small area. At very low frequencies, the magnitude of the response of the passive elastic joint is small and results in a gait with a small area. At mid range frequencies, however, the response of the passive elastic joint has a larger amplitude and a phase shift that results in gaits with larger areas.

4.2 Efficiency in the presence of passive elements

4.2.1 Efficiency for fully actuated swimmers

In drag-dominated environments, a common measure of the cost of any motion executed by a swimmer is the power dissipated into the surrounding fluid while executing the motion. A natural choice of definition for efficiency for these systems is thus

$$\eta_1 = \frac{g_\phi}{P_{avg}T}, \tag{4.16}$$

where g_ϕ is the displacement produced in the x direction over a gait cycle ϕ, P_{avg} is the average power dissipated into the surrounding while executing the gait ϕ, and T is the time taken to execute the gait ϕ. Therefore $P_{avg}T$ is the energy dissipated into the surrounding on executing one cycle of the gait, which can be calculated

using (2.10).

Note that this definition of efficiency is the inverse of the mechanical cost of transport used in [3]. The mechanical cost of transport is a widely used efficiency metric in the robotics community, especially while studying legged locomotion. Gaits which optimize our criterion also optimize Lighthill's efficiency [44], which compares the power dissipated while executing the gait to the power dissipated in rigidly translating the swimmer through the fluid; this measure has an advantage over Lighthill in that it allows for effective comparison between systems with different morphologies and does not require designating a reference shape for rigidly dragging the swimmer. An Explanation for why gaits that optimize our criterion also optimize Lighthill's efficiency is provided in appendix C.

The power term in (4.16) depends on the pacing at which the gait is executed as well as the gait kinematics. Becker et al. showed in [7] that the optimal pacing for drag-dominated systems is the one in which the power dissipated remains constant throughout the cycle, i.e., the pacing at which

$$P(t) = P_{avg}. \tag{4.17}$$

for all t. As detailed in [61], for fully actuated swimmers, this insight lets us restrict our optimization to constant power trajectories and utilize a geometric definition of efficiency,

$$\eta_2 = \frac{g_\phi}{s} \tag{4.18}$$

where s is the pathlength of the gait in the shape space as measured under the

power metric. This measure of efficiency lets us abstract out the pacing computation and relate the efficiency of the gait directly to the geometric parameters of the gait (the surface integral of the CCF and the metric weighted pathlength). As discussed in [59], we can relate the two efficiency measures by

$$\frac{1}{\eta_1} = \frac{1}{\eta_2} + \sigma(\phi) \tag{4.19}$$

where $\sigma(\phi)$ measures the extent to which the system deviates from constant power pacing. It can be represented by squaring the difference between the average and instantaneous rates at which the gait is being followed (measured as s per time), and then integrating over the time during which the gait is being executed,

$$\sigma = \int_0^{\tau_{total}} \left(\frac{s_{\text{total}}}{\tau_{\text{total}}} - \frac{d}{d\tau}(s(\tau))\Big|_{\tau=t} \right)^2 dt, \tag{4.20}$$

where τ_{total} is the time period of the gait, s_{total} is the length of the gait under the metric $\mathcal{M}$, and s is distance traveled along the gait as a function of time corresponding to the given pacing. If the gait is proceeding at constant power, $\frac{s_{\text{total}}}{\tau_{\text{total}}}$ is equal to the rate at which s changes with time, so σ measures the extent to which the pacing lags and leads the optimal pacing.

For any gait ϕ with constant power pacing, $\sigma(\phi) = 0$. The geometric definition of efficiency, η_2 was used to establish the geometric framework in [59, 60].

4.2.2 Efficiency for swimmers with passive elastic joints

In fully actuated swimmers, the transformation of efficiency from the inverse of cost of transport as defined in (4.16) to a more geometric definition in (4.18) was possible because we knew that the optimal pacing keeps the rate of power dissipation constant [7]. In the case of swimmers with passive joints, the response of the passive joint is dictated by the dynamics of the active joint, as illustrated by the Bode plots shown in Fig. 4.4. As a result, there is a unique pacing associated with every gait the passive swimmer can execute. Changing the pacing of the actuated joint changes the response of the passive joint as shown in §4.1 and hence the shape of the gait. Thus to find the most efficient gait for the Purcell swimmer with a passive elastic joint, we have to directly use the definition of efficiency in (4.16).

While we could choose a constant power pacing for all gaits for fully actuated swimmers, in the case of the Purcell and T-link swimmers with a passive elastic joint, every gait that respects the passive dynamics of the elastic joint comes with an inherent pacing. Thus for a given spring stiffness of the passive joint, the gait that maximizes forward velocity comes with a power requirement associated with it. Even if we are capable of giving the system more power, there is no way for the system to utilize that power to go faster. Hence for the swimming systems considered in this paper, there are two meaningful measures for comparing different gaits that lead to different definitions of gait optimality: Gaits can be compared by

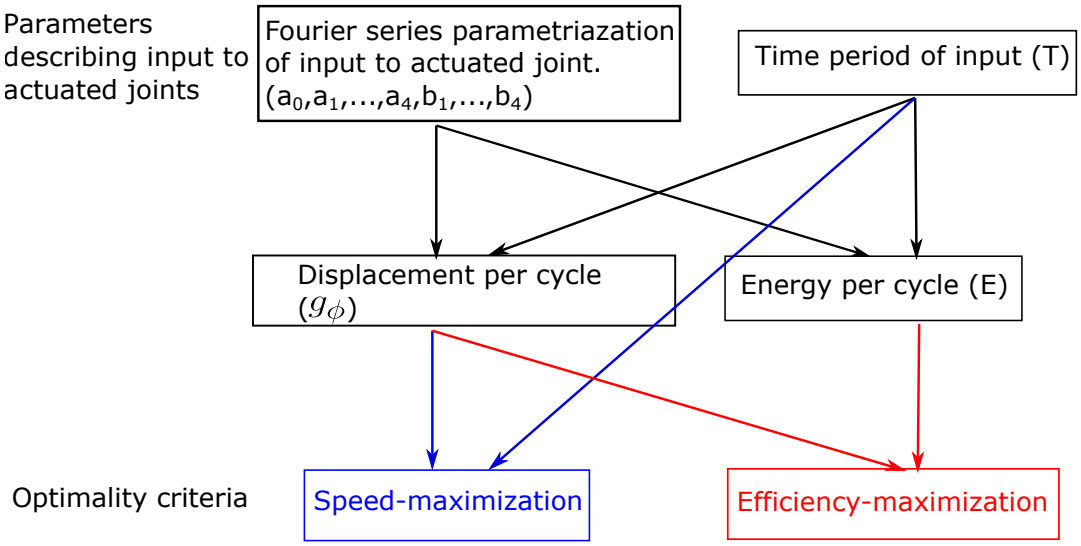

Figure 4.5: A flowchart describing how the parameters describing our input to the actuated joint affect the two optimality criteria in the case of drag-dominated swimmers with a passive joint.

1. Comparing the average speeds they produce ($\eta = \frac{g_\phi}{T}$)
2. Comparing their energetic efficiency (η_1 in (4.16))

4.2.3 Gait parametrization for passive swimmers

We use a truncated Fourier series to parametrize the gait. This choice of parametrization lets us accurately approximate a large family of smooth periodic gaits. The framework introduced in [61] uses a gradient descent algorithm to identify gaits that maximize efficiency as defined in (4.18). During the gradient calculation process outlined in §3.1, it is useful to think of the gait as being parametrized by a

series of waypoints. We can generate these waypoints from the Fourier parametrization. We use the gradients calculated at each of these waypoints to calculate gradients with respect to the Fourier series parametrization.

In the case of swimmers with a passive joint, we let the actuated joint trajectory $\alpha_1(t)$ be given by a fourth order Fourier series,

$$\alpha_1(t) = a_0 + \sum_{i=1}^{4} a_i \cos\left(\frac{2\pi i}{T}t\right) + b_i \sin\left(\frac{2\pi i}{T}t\right) \tag{4.21}$$

Using (4.8), i.e. the transfer function relating the movement of the active and passive joints, we obtain the response of passive joint to $\alpha_1(t)$ as

$$\alpha_2(t) = \mathcal{L}^{-1}(H(s)\mathcal{L}(\alpha_1(t)). \tag{4.22}$$

Using explicit evaluation of the transfer function from (4.10) and (4.12), we can write the steady state response of the passive joint as

$$\alpha_2(t) = \sum_{i=1}^{4} c_i \cos\left(\frac{2\pi i}{T}t\right) + d_i \sin\left(\frac{2\pi i}{T}t\right), \tag{4.23}$$

where c_i and d_i are functions of a_i and b_i and T.

Using this low-order Fourier series parameterization of the gait, we can generate the direct transcription waypoints, calculate the gradient of the objective function at each waypoint (details of this gradient calculation process for speed-maximizing and efficiency-maximizing gaits are presented in §4.3 and §4.4 respectively), then

project these gradients onto the Fourier basis, to obtain gradients with respect to the Fourier series parameters.

If the system were fully actuated, we could move the Fourier series parameters along these calculated gradient directions to obtain the optimal gait. In the case of passive swimmers, the Fourier coefficients of the unactuated shape direction (c_i, d_i) are functions of the Fourier shape coefficients of actuated shape direction (a_i, b_i). Therefore, to find the correct gradient directions for the Fourier coefficients of the actuated shape, we have to account for the change in the unactuated shape direction that a change in the actuated shape direction would produce.

For an objective function f that depends on a_i, b_i, c_i and d_i, we can calculate the total derivatives of f with respect to a_i and b_i as

$$\frac{df}{da_i} = \frac{\partial f}{\partial a_i} + \frac{\partial c_i}{\partial a_i}\frac{\partial f}{\partial c_i} + \frac{\partial d_i}{\partial a_i}\frac{\partial f}{\partial d_i} \tag{4.24}$$

$$\frac{df}{db_i} = \frac{\partial f}{\partial b_i} + \frac{\partial c_i}{\partial b_i}\frac{\partial f}{\partial c_i} + \frac{\partial d_i}{\partial b_i}\frac{\partial f}{\partial d_i}, \tag{4.25}$$

where $\frac{\partial c_i}{\partial a_i}, \frac{\partial c_i}{\partial b_i}, \frac{\partial d_i}{\partial a_i}$ and $\frac{\partial d_i}{\partial b_i}$ are directly taken from the transfer function coefficients (4.15)

$$\frac{\partial c_i}{\partial a_i} = A_2 \tag{4.26}$$

$$\frac{\partial c_i}{\partial b_i} = A_3 \tag{4.27}$$

$$\frac{\partial d_i}{\partial a_i} = -A_2 \tag{4.28}$$

$$\frac{\partial d_i}{\partial b_i} = A_3. \tag{4.29}$$

We use these total derivatives to calculate the correct gradient directions for the Fourier coefficients of the actuated shape variables, which account for the fact that in passive swimmers, a change in the shape of input to the actuated shape variable affects the response of the passive elastic joint.

4.3 Speed-maximizing Gaits

A major goal of this paper is to geometrically identify gaits for the Purcell and T-link swimmers with an elastic joint that will give it the maximum forward velocity. Therefore, the objective function we set out to maximize is $\frac{g_\phi}{T}$, where g_ϕ is the displacement in the x direction produced on executing the gait ϕ, and T is the time period required to execute the gait.

As explained in §4.2.3, we parametrize the actuated shape variable using a low-order Fourier series and obtain the Fourier series parametrization of the resulting motion of the passive joint using the dynamics of the passive elastic joint presented in §4.1. With these Fourier-series parameters, p_f, we can obtain a sequence of waypoints p_i (in our implementation we use 100 waypoints), equally spaced in time, that describe the location of the discretization points in the shape space. As illustrated in this section, we can then calculate the gradient of speed with respect to each of these waypoints, i.e., calculate the effect moving the waypoints would have on the forward speed attained by the swimmer on executing the gait. We then project these gradients onto the Fourier series basis as explained in §4.2.3 to obtain the speed-maximizing gaits using gradient descent.

We start from the basic variational principle that functions reach their extrema when their derivatives go to zero. Given a gait parametrization defined by waypoint parameters p, executed in time T, the maximum-velocity cycle must satisfy the condition that the gradient of forward velocity with respect to the parameters p and T is zero, i.e.

$$\frac{\partial}{\partial p}\left(\frac{g_\phi}{T}\right) = \frac{1}{T}\frac{\partial g_\phi}{\partial p} - \frac{g_\phi}{T^2}\frac{\partial T}{\partial p} \tag{4.30}$$

$$= \frac{1}{T}\frac{\partial g_\phi}{\partial p} = 0, \tag{4.31}$$

and

$$\frac{\partial}{\partial T}\left(\frac{g_\phi}{T}\right) = \frac{1}{T}\frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{T^2} = 0, \tag{4.32}$$

where $\frac{\partial T}{\partial p}$ term in (4.30) can be taken as zero since T and p are two independent variables describing the gait, i.e. the time taken to complete a gait T does not depend on the shape and pacing of the input to the actuated joint described by p.

Once we have the gradient of speed with respect to a direct transcription parametrization p obtained from a Fourier series parametrization p_f, we follow the process outlined in §4.2.3, specifically (4.24) and (4.25), to obtain the gradient of speed with respect to the Fourier series parametrization, $\frac{\partial}{\partial p_f}\left(\frac{g_\phi}{T}\right)$. A graphical depiction of how elements of the Fourier series parametrization affect the gait optimization process is shown in Fig. 4.5.

For suitable seed values p_{f0} and T_0, the maximum-velocity gaits can thus be

obtained by finding the equilibrium of the dynamical system

$$\dot{p}_f = \frac{\partial}{\partial p_f}\Big(\frac{g_\phi}{T}\Big) \tag{4.33}$$

$$\dot{T} = \frac{\partial}{\partial T}\Big(\frac{g_\phi}{T}\Big) = \frac{1}{T}\frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{T^2}. \tag{4.34}$$

In the geometric framework we introduced in [59], we showed that the process of finding efficient gaits for the fully actuated Purcell swimmer is akin to the dynamics of a soap-bubble in which internal pressure and surface tension combine to determine the shape, size and surface concentration of the soap bubble. This is not the case in swimmers with a passive elastic joint. Instead of two independent processes, the optimization process for finding the fastest gait is more unified where (4.33) is the equation that helps obtain the shape of the optimal input to the actuated joint and (4.34) is the equation that helps obtain the optimal pacing of the input.

4.3.1 Shape gradient of the optimal input to the actuated joint.

As discussed in [61], the gradient that affects the shape of the input to the actuated joint, $\frac{dg_\phi}{dp}$, pushes the gait towards maximum displacement cycles. From (5.5), and the fact that variations in p affect ϕ but not the underlying $D\mathbf{A}$ structure, we can reduce this term to

$$\frac{\partial g_\phi}{\partial p} \approx \nabla_p \iint_{\phi_a} (-D\mathbf{A}). \tag{4.35}$$

A powerful geometric principle (the general form of the *Leibniz integral rule* [22]) tells us that *the gradient of an integral with respect to variations of its boundary is equal to the gradient of the boundary with respect to these variations, multiplied by the integrand evaluated along the boundary,* and allows us to rewrite (4.35) in terms of the constraint curvature value and how changes in parameter values move the gait's trajectory through the shape space.

Formally, the multiplication of the gradient of the boundary and the integrand evaluated along the boundary is the *interior product* [1] of the boundary gradient with the integrand,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}). \tag{4.36}$$

In systems with just two shape variables, the interior product reduces to a simple multiplication between the outward component of $\nabla_p \phi$ and the scalar magnitude of the Lie bracket,

$$\nabla_p \iint_{\phi_a} (-D\mathbf{A}) = \oint_{\phi} (\nabla_{p_\perp} \phi)(-D\mathbf{A}). \tag{4.37}$$

This gradient calculation is illustrated in Fig. 3.1. The gradient of the enclosed area with respect to variations in the position of p_i, i.e. $\nabla_{p_i} \phi_a$ in the $e_\parallel$ and $e_\perp$ directions, is the change in triangle's area as p_i moves. Because the triangle's area is always one half base times height (regardless of its pitch or the ratio of its

[1] Not the inner product, see [22] for more details.

sidelengths), this gradient evaluates to

$$\nabla_{p_i}\phi_a = \begin{bmatrix} e_{\|} & e_{\perp} \end{bmatrix} \begin{bmatrix} 0 \\ \ell/2 \end{bmatrix}. \tag{4.38}$$

Note that this term matches the right-hand side of (4.37), with only normal motions of the boundary affecting the enclosed area.

4.3.2 Frequency gradient of the optimal input to the actuated joint

In the case of the fully actuated Purcell swimmer, the shape of the gait, and therefore the displacement produced by executing the gait, are independent of the time taken to execute the gait. This is not true in the case of the Purcell swimmer with a passive elastic joint.

In this subsection, we examine the gradient that guides the optimizer towards the optimal frequency of input to the actuated joint. When the time period required to execute the gait is changed, the shape of the gait changes due to the coupling between the frequency of input to the actuated joint and the response of the passive joint as described in §4.1. Changing the time period T thus changes not only the frequency of the gait cycle but also the displacement produced per cycle.

We use the chain rule to calculate this gradient,

$$\frac{\partial}{\partial T}\left(\frac{g_\phi}{T}\right) = \frac{1}{T}\frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{T^2} \tag{4.39}$$

$$= \frac{1}{T}\left(\frac{\partial g_\phi}{\partial \alpha_1}\frac{\partial \alpha_1}{\partial T} + \frac{\partial g_\phi}{\partial \alpha_2}\frac{\partial \alpha_2}{\partial T}\right) - \frac{g_\phi}{T^2} \tag{4.40}$$

Because α_1 is the actuated shape variable and the shape of the input actuation is independent of the frequency of actuation, $\frac{\partial \alpha_1}{\partial T}$ reduces to zero. Therefore, the gradient of speed with respect to T reduces to

$$\frac{\partial}{\partial T}\left(\frac{g_\phi}{T}\right) = \frac{1}{T}\left(\frac{\partial g_\phi}{\partial \alpha_2}\frac{\partial \alpha_2}{\partial T}\right) - \frac{g_\phi}{T^2}. \tag{4.41}$$

The first term of the right hand side of (4.41) captures the contribution to the velocity of the gait caused by the change in the shape of the gait resulting from a change in T. The second term accounts for the fact that, even without a change in the shape of the gait, an increase in the time required to execute the gait would result in a decrease in the velocity of the gait.

4.3.3 Passive Purcell and T-Link swimmers

We implemented the optimizer described in the §4.3 in Matlab by providing (4.33) and (4.34) as the gradient for the `fmincon` optimizer using the sqp algorithm. The shape of the gait obtained is illustrated in Fig. 4.6(a). Fig. 4.6(b) shows the power

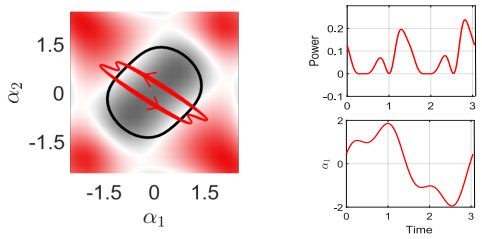

(a) Speed-maximizing gait and pacing

(b) Shape and power of input to the actuated joint of the passive Purcell swimmer

	Passive Purcell Swimmer	Fully Actuated Purcell Swimmer
Passive gait same avg power	8.9×10^{-3}	10.1×10^{-3}
Passive gait same max power	8.9×10^{-3}	18.8×10^{-3}
Active gait same avg power		28.7×10^{-3}
Active gait same max power		53.6×10^{-3}

(c) Maximum speeds achieved by the passive and active Purcell swimmer

Units= $\frac{k}{F_d l_s}$

l_s = length of the swimmer

k = stiffness of the passive joint

F_d = drag force per unit length per velocity

Figure 4.6: Gaits that maximize speed along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated Purcell swimmers at different power levels.

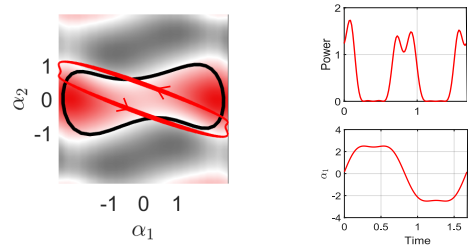

(a) Speed-maximizing gait and pacing

(b) Shape and power of input to the actuated joint of the passive T-link swimmer

	Passive T-link Swimmer	Fully Actuated T-link Swimmer
Passive gait same avg power	1.4×10^{-2}	1.8×10^{-2}
Passive gait same maxpower	1.4×10^{-2}	3.0×10^{-2}
Active gait same avg power		5.3×10^{-2}
Active gait same max power		9.1×10^{-2}

Units of speed: Body lengths (l) per second (s)

(c) Maximum speeds achieved by the passive and active T-link swimmer

Units= $\frac{k}{F_d l_s}$

l_s = length of the swimmer

k = stiffness of the passive joint

F_d = drag force per unit length per velocity

Figure 4.7: Gaits that maximize speed along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated T-link swimmers at different power levels.

input to the actuated joint over the cycle. Fig. 4.6(c) shows a comparison of the speeds achievable by the Passive and fully actuated Purcell swimmers at different power levels. Fig. 4.7 shows the same results for a passive T-link swimmer.

The transfer function relating the response of the passive joint to oscillations of the input joint is given by (4.9), $H(s) = \dfrac{s}{C_{eq}s + K_{eq}}$. Therefore a change in the value of the spring stiffness does not affect the fundamental shape and nature of the response shown by the Bode plot in Fig. 4.4, but it does shift the entire bode plot to the left or right along the frequency axis. Thus an increase in spring stiffness shifts the Bode plot to the right, which results in the shape of the speed maximizing gait remaining the same, but the time period required to complete the gait decreases, leading to faster speeds.

4.4 Energy-Efficient Gaits

In this section, we describe the gradient calculations involved in identifying the gait that maximizes the efficiency of the swimmers. The objective function we set out to maximize is

$$\eta = \frac{g_\phi}{E}, \tag{4.42}$$

where g_ϕ is the displacement produced on executing the gait ϕ and E is the total energy expended by the robot executing the gait, i.e.

$$E = P_{avg}T \tag{4.43}$$

As explained in §4.2.3, we parametrize the actuated shape variable using a low-order Fourier series, p_f, and obtain the Fourier series parametrization of the resulting motion of the passive joint using the dynamics of the passive elastic joint presented in §4.1. With these Fourier-series parameters, we can obtain a sequence of waypoints p_i equally spaced in time, that explicitly define the location of the discretization points in the shape space. These waypoints form a direct-transcription parametrization of the gait p, obtained from the Fourier series parametrization p_f. In this section, we present the calculation of the gradient of efficiency at each of these points with respect to p. We then project these gradients onto the Fourier series basis as explained in §4.2.3 to obtain the efficiency maximizing gaits using gradient descent.

The maximum-efficiency cycle must satisfy the condition that the gradient of efficiency with respect to the parameters, p and T is zero, i.e.

$$\begin{aligned} \frac{\partial}{\partial p}\Big(\frac{g_\phi}{E}\Big) &= \frac{1}{E}\frac{\partial g_\phi}{\partial p} - \frac{g_\phi}{E^2}\frac{\partial E}{\partial p} \\ &= 0 \end{aligned} \tag{4.44}$$

and

$$\begin{aligned} \frac{\partial}{\partial T}\Big(\frac{g_\phi}{E}\Big) &= \frac{1}{E}\frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{E^2}\frac{\partial E}{\partial T} \\ &= 0 \end{aligned} \tag{4.45}$$

Once we have the gradient of efficiency with respect to a direct transcription

parametrization p obtained from a Fourier series parametrization p_f, we follow the process outlined in §4.2.3 to obtain the gradient of efficiency with respect to the Fourier series parametrization, $\frac{\partial}{\partial p_f}\left(\frac{g_\phi}{E}\right)$.

Thus for suitable seed values of p_{f0} and T_0, the maximum efficiency gaits can be obtained by finding the equilibrium of the dynamical system,

$$\dot{p}_f = \frac{\partial}{\partial p_f}\left(\frac{g_\phi}{E}\right) \tag{4.46}$$

$$\dot{T} = \frac{1}{E}\frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{E^2}\frac{\partial E}{\partial T} \tag{4.47}$$

Thus the process of finding the most efficient gait is the result of a unified process where (4.46) is the equation that helps find the shape of the input to the actuated joint and (4.47) is the equation that helps find the optimal pacing of the input. Note that as in §4.3, the two equations do not operate independently, and the gradient of shape depends on the time period T, and the gradient of the time period depends on the shape of the gait.

4.4.1 Shape gradient of the optimal input to the actuated joint

The shape of the optimal input to the actuated joint is affected by two gradients, $\frac{\partial g_\phi}{\partial p}$and $\frac{\partial E}{\partial p}$ (4.46). The details of how $\frac{\partial g_\phi}{\partial p}$ is calculated are explained in §4.3. Whereas $\frac{\partial g_\phi}{\partial p}$ pushes the gait towards maximum displacement cycles, E is a measure

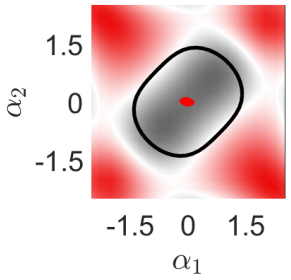

(a) Efficiency-maximizing gait and pacing

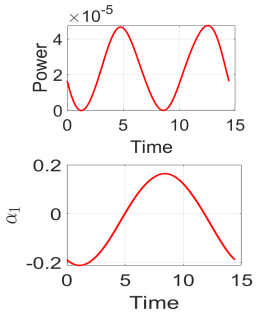

(b) Shape and power of input to the actuated joint of the passive Purcell swimmer

	Passive Purcell Swimmer	Fully Actuated Purcell Swimmer
Passive gait same speed	1.33	1.99
Active gait same speed		221.62

Units= $\frac{l_s}{k}$

l_s = length of the swimmer

k = stiffness of the passive joint

(c) Maximum efficiencies achieved by the passive and active Purcell swimmer

Figure 4.8: Gaits that maximize efficiency along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated Purcell swimmers moving forward at the same speed as the passive swimmer. Note that for the passive Purcell swimmer, the optimizer stops because reducing the frequency further or making the gait smaller does not provide any meaningful increase in efficiency. This observation is line with the results in [56].

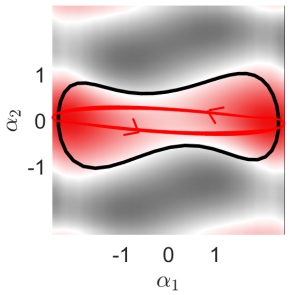

(a) Efficiency-maximizing gait and pacing

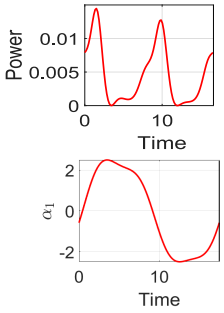

(b) Shape and power of input to the actuated joint of the passive T-link swimmer

	Passive T-link Swimmer	Fully Actuated T-link Swimmer
Passive gait same speed	0.32	0.42
Active gait same speed		3.12

Units= $\frac{l_s}{k}$

l_s = length of the swimmer

k = stiffness of the passive joint

(c) Maximum efficiencies achieved by the passive and active T-link swimmer

Figure 4.9: Gaits that maximize efficiency along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated T-link swimmers moving forward at the same speed as the passive swimmer.

of the cost required to execute the gait and $\frac{\partial E}{\partial p}$ pushes the gait towards low cost shapes. At the most efficient gait, these two opposing gradients cancel each other out, and we get an equilibrium for the gait optimization process.

Over a gait cycle, no energy is stored in the spring. Hence we can calculate the energy expended, P, while executing a gait by integrating the power flow through the actuated joint (α_1).

$$E = \int_0^T \dot{\alpha}_1(t)^T \tau_1 dt \tag{4.48}$$

$$= \int_0^T \dot{\alpha}_1(t)^T \mathcal{M}_1(t)\dot{\alpha}(t)dt \tag{4.49}$$

where $\mathcal{M}_1(t)$ is the first row of the power metric $\mathcal{M}(t)$. The gradient of cost with respect to the shape of the gait, $\frac{\partial E}{\partial p}$, is calculated by

$$\frac{\partial E}{\partial p} = \frac{\partial}{\partial p}\int_0^T \dot{\alpha}_1(t)^T \mathcal{M}_1(t)\dot{\alpha}(t)dt \tag{4.50}$$

$$= \int_0^T \Big(\frac{\partial \dot{\alpha}_1}{\partial p}\mathcal{M}_1\dot{\alpha} + \tag{4.51}$$

$$\dot{\alpha}_1^T \mathcal{M}_1 \frac{\partial \dot{\alpha}}{\partial p} + \dot{\alpha}_1^T \frac{\partial \mathcal{M}_1}{\partial p}\dot{\alpha}\Big)dt \tag{4.52}$$

4.4.2 Frequency gradient of the optimal input to the actuated joint

The equation that governs the optimization process for finding the time period of the most efficient gait is described by (4.47). The term $\frac{\partial g_\phi}{\partial T}$ is calculated as described in §4.3.2. The second gradient in the right hand side of (4.47) is calculated

as

$$\frac{dE}{dT} = \frac{\partial E}{\partial \alpha_1}\frac{\partial \alpha_1}{\partial T} + \frac{\partial E}{\partial \alpha_2}\frac{\partial \alpha_2}{\partial T} + \frac{\partial E}{\partial T}. \tag{4.53}$$

Because α_1 is the actuated shape variable and the shape of the input actuation is independent of the frequency of actuation, $\frac{\partial \alpha_1}{\partial T}$ reduces to zero. Therefore the gradient of energy with respect to period reduces to

$$\frac{dE}{dT} = \frac{\partial E}{\partial \alpha_2}\frac{\partial \alpha_2}{\partial T} + \frac{\partial E}{\partial T}. \tag{4.54}$$

The first term accounts for the fact that a change in frequency would change the response of the passive joint α_2 resulting in a change in the shape of the gait, and hence a change in the power dissipated while executing the gait. The second term accounts for the fact that even if the shape of the gait remains unchanged, a change in the frequency of input to the actuated joint will change the time required to execute the gait and hence would change the power dissipated while executing the gait.

4.4.3 Passive Purcell and T-Link swimmers

We implemented the optimizer described in the §4.4 in Matlab by providing the gradients of efficiency with respect to shape and time period, calculated using (4.46) and (4.47) respectively, to the `fmincon` optimizer using the sqp algorithm. The shape of the gait obtained is illustrated in Fig. 4.8(a) for a Purcell swimmer.

Fig. 4.8(b) shows the power input to the actuated joint over the cycle. Fig. 4.6(c) shows a comparison of the efficiencies achievable by the Passive and fully actuated Purcell swimmers when the forward speed for all the systems is fixed to be equal to the forward speed achieved by the passive swimmer when executing its maximum efficiency cycle. Fig. 4.9 shows the same results for a passive T-link swimmer. Note that for the Purcell swimmer, the optimizer stops because reducing the frequency further or making the gait smaller does not provide any meaningful increase in efficiency. This observation is line with the results from [56], where the efficiency was found to asymptotically approach a maximum value as frequency of gait oscillations approached zero.

As discussed in §4.3.C, change in spring stiffness does not affect the shape and nature of the response of the passive joint, but shifts the bode plot shown in Fig. 4.4 to the left or right along the frequency axis. An increase in spring stiffness shifts the bode plot to the right, which results in the shape of the efficiency maximizing gait cycle remaining the same, but the time taken to execute the gait will decrease. The energy required to execute a gait is inversely proportional to the time taken to execute the gait. Therefore, more energy is required to execute a gait faster. Thus, increasing spring stiffness will result in an overall decrease in the efficiency of swimming.

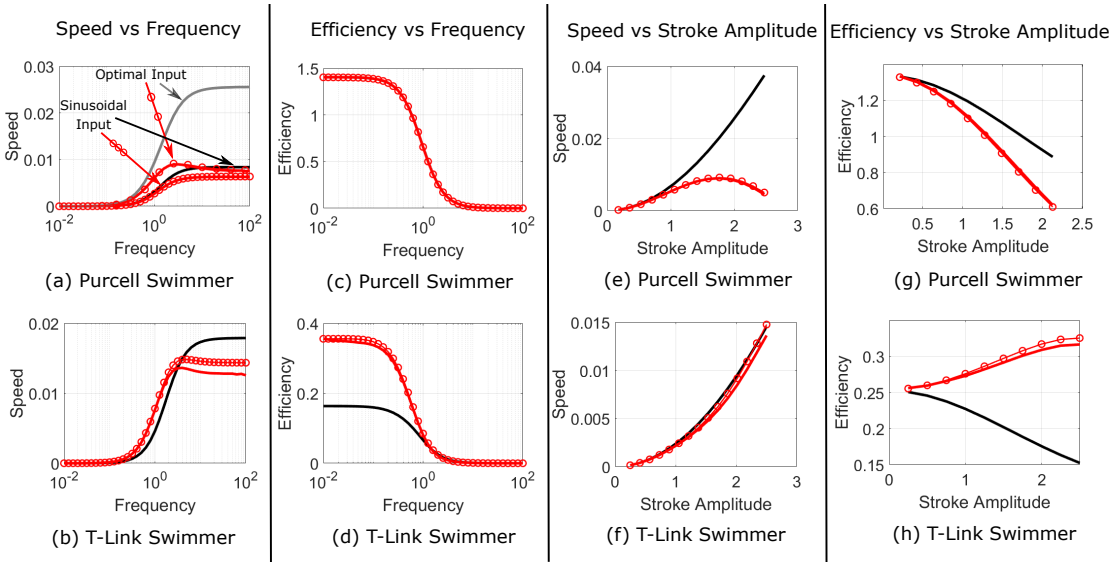

Figure 4.10: In all the subfigures, the solid red lines and the red circles show the speeds and efficiencies predicted by numerical simulation and integral of CCF respectively. The solid black lines show the speeds and efficiencies predicted by the constant-CCF assumption used in [56] in the link-attached coordinates. The first column of figures illustrate the speed of the (a) Purcell swimmer as a function of actuation frequency when the input to the controlled joint is a sinusoidal oscillation of unit amplitude and when the input to the controlled joint has the optimal shape and (b) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §4.3. In (a), the grey line shows the speed predicted by the constant-CCF assumption for a sinusoidal input. The small amplitude perturbation analysis in [56] predicts speed to be a monotonically increasing function of frequency for all inputs. The speed is a monotonically increasing function of frequency for the sinusoidal input. However, it is not monotonically increasing for the optimal input contrary to the prediction from the small amplitude perturbation analysis in [56]. The second column of figures illustrate the efficiency of the (c) Purcell swimmer and (d) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §4.4. Figures (e)-(h) illustrate the speed and efficiency of the Purcell swimmer and T-link swimmer as a function of the magnitude of optimally shaped input actuation obtained from §4.4.

4.5 Comparison with previous work

Passive swimmers have been previously studied, e.g., in [15, 38, 42, 49, 56]. Of these works, [42, 49] and [56] are most relevant to this paper, as they discuss the motion of swimmers with a harmonically-driven active joint and a passive joint. The T-link swimmer used in this paper was introduced in [42] as a simplified model to study the swimming dynamics of *Schistosoma mansoni*. The analysis in [56] is particularly relevant, as it applies perturbation theory to investigate the motion of the Purcell swimmer.

The approximation used in the perturbation analysis in [56] is equivalent to assuming that the displacement produced by executing a gait is equal to the area enclosed by the gait in the shape space multiplied by the constraint curvature value at the center of the gait,

$$g_\phi \approx \iint_{\phi_a} D(-\mathbf{A})|_{\mathbf{0}} = \phi_a \cdot D(-\mathbf{A})|_{\mathbf{0}}. \tag{4.55}$$

This approximation has been used in several works from the geometric mechanics community to identify useful shape oscillations that resulted in useful net displacements e.g., [50, 52].

Using the approximation in (4.55), the authors of [56] concluded that for harmonic inputs, the speed of the swimmer monotonically increases with frequency and asymptotically approaches a maximum value as the actuation frequency approaches ∞. They similarly concluded that the efficiency of the swimmer asymptotically approaches a maximum value as the actuation frequency approaches zero.

A drawback of this approximation is that the accuracy falls steeply with increasing gait size, as a larger gait would mean larger variations in the value of CCF inside the region bounded by the gait. Our analysis improves on this approximation in three ways: first, our use of a body-averaged frame instead of the link-attached frame in [56] "flattens" some of the nonlinearity in the system, expanding the domain of gait amplitude for which the perturbation analysis gives accurate results. Second, our use of the full integral of constraint curvature over the area enclosed by the gait,

$$g_\phi \approx \iint_{\phi_a} D(-\mathbf{A}), \tag{4.56}$$

absorbs much of the remaining nonlinearity. Third, we use (4.56) only to calculate gradients, and numerically evaluate the value of g_ϕ at each step of the gradient descent process described in §4.3 and §4.4 to avoid compounding errors from any residual nonlinearity

Fig. 4.10(a) and Fig. 4.10(b) show how the effect on swimming speed from changing the frequency of the input stroke for Purcell and T-link swimmers respectively. The solid red lines and the red circles show the speeds predicted by numerical simulation and area integral of curvature respectively. The solid black and grey lines show the speeds predicted by the constant-CCF assumption used in [56] in the link-attached coordinates for optimally shaped and sinusoidal inputs respectively. The speeds predicted by the constant-CCF assumption are higher than the actual speeds obtained by numerical simulations, and our integral of

CCF is a good approximation of the ground-truth simulation.

We can also see from Fig. 4.10(a) that the velocity obtained by the Purcell swimmer does not monotonically increase with frequency for all inputs to the actuated joint. This results in the speed-maximizing gaits found in §4.3 having an optimal frequency associated with them, rather than exhibiting a monotonic increase in speed with frequency.

In the case of efficiency-maximizing gaits, for the optimal gait shape obtained in §4.4, the efficiency does asymptotically approach a maximum value as shown in Fig. 4.10(c) and Fig. 4.10(d) as the frequency approaches zero, but the value is different from the maximum efficiency predicted by applying the small perturbation analysis from [56] to the T-link swimmer, showing that small perturbation analysis does not completely characterize optimal performance. In the case of the Purcell swimmer, the efficiency-maximizing gait is small enough for the perturbation analysis to yield accurate results.

Fig. 4.10(e)-(h) show how the constant-CCF assumption can introduce errors in identifying optimal actuation shape. Fig. 4.10(e) and Fig. 4.10(f) show the effect on the swimming speed from changing the size of the input stroke (the reference input is the optimal input obtained in §4.3). Fig. 4.10(g) and Fig. 4.10(h) show the effect on the swimming efficiency from changing the size of the input stroke (the reference input is the optimal input obtained in §4.4). The solid red lines and the red circles show the speeds and efficiencies predicted by numerical simulation and integral of CCF respectively. The solid black lines show the speeds and efficiencies predicted by the constant-CCF assumption used in [56] in the link-attached coordinates.

We can see that in the case of the Purcell swimmer, the constant-CCF assumption used in [56] incorrectly predicts a monotonous increase in speed with an increase in the amplitude of the input to the actuated joint. In the case of the T-link swimmer, the constant-CCF assumption used in [56] incorrectly predicts an increase in efficiency as we shrink the optimal gait. The efficiency would go down if we shrink the optimal gait, because the CCF value for T-link swimmer is higher at the edges than at the center of the shape space. When we shrink the gait, it loses these regions of high value, leading to a decrease in efficiency, which is not captured by the constant-CCF assumption.

The T-link swimmer was first introduced in [42]. The analysis in this paper agrees with the most relevant results from [42], which are that when the actuated joint is driven by a simple harmonic input:

1. There exists a linear relationship between the speed-maximizing value of spring stiffness and frequency of actuation.
2. The average swimming speed increases monotonically with the amplitude of actuation from $\frac{\pi}{2}$ radians to $\frac{11\pi}{9}$ radians.

In §4.3.3, we noted that an increase or decrease in spring stiffness shifts the Bode plot of the response of the passive joint to the right or to the left without changing the shape of the Bode plot resulting in frequency of the speed-maximizing input being linearly related to the stiffness of the passive joint.

In Fig. 4.7(a), we can see that the CCF value for T-link swimmer is higher at the edges than at the center of the shape space. Thus an increase in the amplitude

of actuation would enclose more of the high value region leading to an increase in speed.

4.6 Conclusions

In this paper, we have identified the geometric structure of optimal gaits for viscous swimmers with passive elastic joints by combining the constraint-curvature analysis in [61] with frequency-response models for the steady state motion of driven oscillators. We use this structure to identify both speed-maximizing and efficiency-maximizing gaits. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

For the fully actuated swimmers, the maximum forward speed achievable is only restricted by the maximum power we are able to supply the joints, but for the swimmers with the passive elastic joint, even with more powerful actuators, there is a theoretical maximum forward speed the system can achieve dictated by the stiffness of the passive joint.

The important factor that makes the performance of the fully actuated swimmers superior to that of the swimmers with the passive elastic joint in terms of energy efficiency is the fact that not only can the fully actuated swimmers execute any gait the passive swimmers can execute, they can do so at a pacing that is just as good or better than the pacing dictated by the dynamics of the passive joint, which is the only pacing at which the passive swimmers can execute the gait.

This raises the question of what benefits, e.g., simplicity of construction, does having a passive elastic member give to biological organisms that locomote in a low Reynolds number fluid? Most biological organisms have tails that resemble an elastic filament. The propulsive and flexive dynamics of such filaments have been well studied [13, 25, 69, 70]. Artificial microscopic swimmers with elastic filaments have been proposed based on this body of work [19, 20]. An interesting line of future work would involve investigating the tradeoff between elastic element inefficiencies and structural complexity of being fully actuated.

This work is the first step towards expanding the applicability of the geometric framework presented in [59–61] to systems where underactuated shape parameters play a role in the dynamics of the system. In the case of the passive Purcell swimmer, assuming the torque required to affect a desired shape change did not depend on the current shape of the swimmer did not introduce significant errors in the predicition of the limit cycle corresponding to inputs to the actuated joint as shown in Fig. 4.1. This might not always be the case in all the swimmers we consider. A future line of research would be to improve our frequency domain analysis by using non-linear perturbation theory to obtain more accurate predictions of limit cycles.

Chapter 5: Magnus Expansion and Application to Gait Analysis

Many kinematic locomoting systems can be modeled as time-varying systems with a linear relationship between the rate of internal deformation and rate of movement through the world. Several efforts in the geometric mechanics community [4, 46, 50, 52, 53, 64, 68] (including our own [32, 33]), have aimed at simplification of the BCH expansion in attempts to use *curvature* of the system constraints (a measure of how "non-canceling" the system dynamics are) to understand which gaits produce useful displacements. This resulted in works that showed Lie brackets of the system dynamics can be used to approximately identify shape oscillations that produce useful net displacements [50, 52], and that the scale of motions for which the curvature methods provide accurate information can be significantly extended by pre-optimizing the choice of system coordinates [34]. By working in these new coordinates, we have gained insight into the optimal gaits of swimmers and crawlers in low- and high-viscosity fluids [28, 32, 33], and in granular media [27]. We used these insights to build a variational framework to identify optimal gaits for drag dominated kinematic systems [59, 60].

Since all these works base their analysis on the total Lie bracket formula obtained by restricting the BCH expansion to the first few terms, it is crucial that the truncation be a good approximation of the full BCH expansion. The aim of the paper is to understand the factors affecting the accuracy of the total Lie bracket

formula. We show that the major factors affecting the accuracy are the magnitude of the connection vector field and the size of the gait. Because the optimal coordinate choice minimizes the magnitude of the connection vector field, it improves the accuracy of the truncated Magnus series and hence the accuracy of the total Lie bracket formula.

The rest of the paper is organized as follows: In §5.1 we present a derivation of the total Lie bracket formula used to identify optimal gaits. We show that the accuracy of the total Lie bracket formula hinges on two sets of assumptions. In §5.2, we relate the accuracy of these assumptions to the system dynamics and the physical parameters of the gait. In §5.3, we present a heuristic for selecting a starting point that minimizes the leading order error associated with the total Lie bracket formula. We then demonstrate the impact of operating in optimal coordinates on the total Lie bracket accuracy for our example systems.

5.1 Gaits and BCH formula

5.1.1 Gaits

Many locomoting systems, including our example systems move via gaits: cyclic changes in shape that produce characteristic net displacements. We now recall some of the conventions associated with gaits introduced in [34].

***Definition* 1** *A gait ϕ with a time period of T is a periodic mapping from time $\mathbb{R}$ to the shape space R, i.e. $\phi : \mathbb{R} \to R$, such that $\phi(t) = \phi(t + nT)$, $\forall n \in \mathbb{Z}$ and*

$\forall t \in \mathbb{R}$.

Therefore the set of all possible gaits is given by,

$$\Phi = \{\phi : \mathbb{R} \to M | \phi(t) = \phi(t+T), \forall t \in \mathbb{R}, \forall n \in \mathbb{Z}\} \tag{5.1}$$

Definition 2 *The image family $\overline{\phi}$ of a gait ϕ is the set of all gaits which share its image (i.e. trace out the same closed curve) in R,*

$$\overline{\phi} = \{\psi \in \Phi | Im(\psi) = Im(\phi)\} \tag{5.2}$$

Note that since we are dealing with kinematic systems, the time period and pacing taken to execute a gait does not affect the displacement produced by the gait, hence from hereon, we will assume that all gaits we consider have a constant time period, T.

The net displacement produced by one cycle of a gait ϕ with a time period T is given by:

$$g_\phi = (g(0))^{-1} g(T) = \int_0^T g(t)\mathbf{A}(r(t))\dot{r}(t)dt \tag{5.3}$$

where $r(t) = \phi(t)$. If we assume that the system starts out at the identity element of the position space G, then $g_\phi = g(T)$.

Several efforts in the geometric mechanics community [4, 46, 50, 52, 53, 64, 68] (including our own [32, 33]), have aimed at using the Baker-Campbell-Hausdorff(BCH)

formula in attempts to use *curvature* of the system constraints (a measure of how "non-canceling" the system dynamics are) to understand which gaits produce useful displacements.

The core principle in these works is that because the net displacement g_ϕ over a gait cycle ϕ is the line integral of (2.1) along ϕ, the displacement can be approximated by an area integral of the curvature $D\mathbf{A}$ of the local connection (its total Lie bracket [34]) over a surface ϕ_a bounded by the cycle:

$$g_\phi = \int_0^T -g(t)\mathbf{A}(r(t))\dot{r}(t)dt = \oint_\phi -g\mathbf{A}(r)dr \tag{5.4}$$

$$\approx \exp\Big(\iint_{\phi_a} \underbrace{-\mathbf{dA} + \textstyle\sum\big[\mathbf{A}_i, \mathbf{A}_{j>i}\big]}_{D\mathbf{A}\text{ (total Lie bracket)}}\Big), \tag{5.5}$$

where $\mathbf{dA}$, the exterior derivative of the local connection (its generalized row-wise curl), measures how changes in $\mathbf{A}$ across the shape space prevent the net induced motions from canceling out over a cycle, and the local Lie bracket $\sum\big[\mathbf{A}_i, \mathbf{A}_{j>i}\big]$ measures how translations and rotations in the induced motions couple into "parallel parking" effects that contribute to the net displacement.

In this chapter, our goal is to understand the factors affecting the quality of this approximation in (5.5). To do this we first present a derivation of (5.5) and relate the quality of the approximation to the characteristics of the system such as the magnitude of the vector fields and characteristics of the gait, such as the circumference and area enclosed on the shape space.

5.1.2 Total Lie Bracket

In this section we present the derivation of the total Lie bracket formula (5.5). We show that the accuracy of (5.5) is dependent on two sets of approximations.

The kinematic locomotion model used to describe the systems in this chapter is

$$\dot{g}(t) = g(t)\mathbf{A}(t). \tag{5.6}$$

Therefore from linear systems theory, we can write the infinitesimal time evolution of the system as

$$g(t + \delta t) = g(t)\exp(\mathbf{A}(t)dt), \tag{5.7}$$

and the complete time evolution as

$$g(T) = g(0)\prod_{t=0}^{t=T}\exp(\mathbf{A}(t)dt) \tag{5.8}$$

where $\prod$ is the product-integral, which is the continuous counterpart to the product of a sequence, much in the same way that the standard integral is the continuous counterpart to summation:

$$\prod_{a}^{b}(I + v(x)dx) = \lim_{n\to\infty}\prod_{k=1}^{n}(I + v(x_k)\delta x) \tag{5.9}$$

where $\delta x = \frac{b-a}{n}$ and $x_k = (a + (k-1)\delta x)$.

For all time time t, $\mathbf{A}(t)$ is an element of the noncommutative Lie algebra $se(2)$. Therefore the product-integral in (5.8) is a product of exponentials of non-

commutative Lie algebra elements. The Baker-Campbell-Hausdorff formula tells us how to evaluate the product of exponentials of two noncommutative Lie algebra elements. If X and Y are two noncommutative Lie algebra elements, the BCH formula says that

$$\begin{aligned}\exp(X)\exp(Y) &= \exp\Big(X+Y+\frac{1}{2}[X,Y]+\frac{1}{12}([X,[X,Y]]+[Y,[Y,X]]) \\ &\quad -\frac{1}{24}([Y,[X,[X,Y]]) \\ &\quad -\frac{1}{720}([Y,[Y,[Y,[Y,X]]]]+[X,[X,[X,[X,Y]]]]) \\ &\quad +\cdots\Big) \end{aligned} \tag{5.10}$$

Thus if X and Y were to commute with each other, i.e. if $[X,Y]=0$, then,

$$\exp(X)\exp(Y)=\exp(X+Y) \tag{5.11}$$

Using this insight, if we divide the gait into four temporally equal parts and assume local commutativity in each of the four section, we get from (5.8),

$$\begin{aligned} g(T) &= g(0)\prod_{t=0}^{t=\frac{T}{4}}\exp(\mathbf{A}(t)dt)\prod_{t=\frac{T}{4}}^{t=\frac{T}{2}}\exp(\mathbf{A}(t)dt)\prod_{t=\frac{T}{2}}^{t=\frac{3T}{4}}\exp(\mathbf{A}(t)dt)\prod_{t=\frac{3T}{4}}^{t=T}\exp(\mathbf{A}(t)dt) \\ &\approx g(0)\exp(a)\exp(b)\exp(c)\exp(d) \end{aligned} \tag{5.12}$$

where

$$a = \exp\left(\int_0^{\frac{T}{4}} dt_1 \mathbf{A}(t_1)\right) \tag{5.13}$$

$$b = \exp\left(\int_{\frac{T}{4}}^{\frac{T}{2}} dt_1 \mathbf{A}(t_1)\right) \tag{5.14}$$

$$c = \exp\left(\int_{\frac{T}{2}}^{\frac{3T}{4}} dt_1 \mathbf{A}(t_1)\right) \tag{5.15}$$

$$d = \exp\left(\int_{\frac{3T}{4}}^{T} dt_1 \mathbf{A}(t_1)\right) \tag{5.16}$$

Thus our first assumption of local commutativity, made in in (5.12), lets us describe the displacement resulting from a gait as

$$g_\phi = g^{-1}(0)g(T) \approx \exp(a)\exp(b)\exp(c)\exp(d) \tag{5.17}$$

where a, b, c and d are elements of the noncommutative Lie algebra $se(2)$ which encode the contribution to the final displacement from each quarter of the gait.

We can use the BCH formula to simplify the expression for total displacement further,

$$
\begin{aligned}
g_\phi &\approx \exp(a)\exp(b)\exp(c)\exp(d) && (5.18)\\
&= \exp\big(a+b+c+d+\frac{1}{2}([a,b]+[c,d]) \\
&\quad +\frac{1}{2}([a,c]+[a,d]+[b,c]+[b,d]) \\
&\quad +\frac{1}{12}([a-b,[a,b]]+[c-d,[c,d]]) \\
&\quad +\frac{1}{4}([a+b,[c,d]]-[c+d,[a,b]]+\frac{1}{8}([[a,b],[c,d]]) \\
&\quad +\cdots\big). && (5.19)
\end{aligned}
$$

To relate the displacement resulting from each quarter of the gait to the magnitude of the connection vector fields and the size of the gait we define new variables α, β, δ and Δ related to a, b, c, d as:

$$
\begin{aligned}
a &= \alpha - \delta/2 && (5.20)\\
b &= \beta + \Delta/2 && (5.21)\\
c &= -(\alpha+\delta/2) = -\alpha - \delta/2 && (5.22)\\
d &= -(\beta-\Delta/2) = -\beta + \Delta/2, && (5.23)
\end{aligned}
$$

In appendix D, we show that assuming a linearly varying connection vector field

over the shape space, α, β, δ and Δ are related to the connection vector field by:

$$\alpha = 2L\mathbf{A}_1|_{(0,0)} \tag{5.24}$$

$$\beta = 2L\mathbf{A}_2|_{(0,0)} \tag{5.25}$$

$$\delta = 4L^2\frac{\partial \mathbf{A}_1}{\partial \alpha_2} \tag{5.26}$$

$$\Delta = 4L^2\frac{\partial \mathbf{A}_2}{\partial \alpha_1}. \tag{5.27}$$

Substituting these variables, (5.20)- (5.23), into the expression for total displacement (5.19), expanding and reducing the Lie brackets, we get,

$$\exp(a)\exp(b)\exp(c)\exp(d) = \exp(-\delta + \Delta + [\alpha, \beta] + \frac{1}{4}[-\delta, \Delta]$$
$$+\frac{1}{2}[(\alpha + \beta + \frac{1}{4}([\alpha, \Delta] + [\beta, \delta])), (-\delta + \Delta + [\alpha, \beta] + \frac{1}{4}[-\delta, \Delta])] + \cdots\big) \tag{5.28}$$

To understand what these terms mean in terms of the gait and connection vector fields of the system, we add equations (5.20)-(5.23) to obtain

$$a + b + c + d = -\delta + \Delta \tag{5.29}$$

The above equation tells us that the sum of the first two terms of the BCH formula in (5.28) is equal to the the sum of the contribution to the displacement from each quarter of the gait. Adding the local commutativity approximation equations (5.13)-(5.16), we can show that the sum of the contribution to the displacement from each quarter of the gait is equal to the area integral of the curl of

the connection vectorfield over the region enclosed by the gait:

$$a + b + c + d = \int_0^T dt_1 \mathbf{A}(t_1) = \iint_\phi d\mathbf{A} \tag{5.30}$$

Therefore, the first two terms of the BCH formula in (5.28) equal the area integral of the curl of the connection vectorfield over the region enclosed by the gait:

$$-\delta + \Delta = a + b + c + d = \int_0^T dt_1 \mathbf{A}(t_1) = \iint_\phi d\mathbf{A}. \tag{5.31}$$

If we approximate $\mathbf{A}$ as a linearly varying vector field over the shape space, i.e., $\frac{\partial \mathbf{A}_i}{\partial r_j}$ is constant throughout the shape space $\forall i, j$, and assume a square gait of perimeter l, the first two terms of the BCH formula in (5.28), $-\delta + \Delta$, and the third term, $[\alpha, \beta]$, are of the order $O(l^2)$, and the other terms are $O(l^3)$ or higher. Details of these calculations are presented in appendix D. We also get,

$$[\alpha, \beta] = \iint_\phi [\mathbf{A}_1, \mathbf{A}_2] \tag{5.32}$$

where $\mathbf{A}_1$ is the first column of $\mathbf{A}$, and $\mathbf{A}_2$ is the second column of $\mathbf{A}$. Note that, although we restrict our analysis to square shaped gaits, we can always remap any simple closed loop in the shape space onto a square through a change of coordinates for the shape space.

In our second approximation we assume all terms of order $O(l^3)$ and higher are

small enough to be ignored, i.e.

$$\begin{aligned} \exp(a)\exp(b)\exp(c)\exp(d) &\approx \exp(-\delta + \Delta + [\alpha, \beta]) && (5.33) \\ &= \exp(\iint_{\phi} d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]). && (5.34) \end{aligned}$$

Therefore,

$$g_\phi \approx \exp(a)\exp(b)\exp(c)\exp(d) \approx \exp(\iint_{\phi} d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]) \quad (5.35)$$

We can now see that the accuracy of our approximation of the displacement produced on executing a gait by the surface integral of the total Lie bracket depends on the accuracy of two assumptions. The first set of assumptions are in (5.13)-(5.16) where we assume local commutativity over each quarter of the gait and the second assumption is in (5.34), where we assume all terms of $O(l^3)$ and higher from the BCH formula are small enough to be ignored (l is the perimeter of the gait).

5.2 Accuracy of Total Lie Bracket

In the previous section, we showed that the accuracy of approximating the displacement produced by a gait using the surface integral of the total Lie bracket depends on two sets of assumptions: local commutativity assumption in (5.12) and restricting ourselves to second order terms of the BCH formula in (5.34). In this section, we relate the accuracy of the two sets of assumptions to the magnitude of the connection vector field and physical parameters of the gait such as its perimeter

and area enclosed.

5.2.1 Accuracy of Approximation I

The local commutativity assumption in (5.12) assumes commutativity in each quarter of the gait. This lets us evaluate the contribution from each quarter of the gait as

$$g^{-1}(0)g(\tfrac{T}{4}) = \prod_{t=0}^{T/4} \exp(\mathbf{A}(t)dt) \approx \exp\left(\int_0^{T/4} dt_1 \mathbf{A}(t_1)\right) \tag{5.36}$$

$$g^{-1}(\tfrac{T}{4})g(\tfrac{T}{2}) = \prod_{T/4}^{T/2} \exp(\mathbf{A}(t)dt) \approx \exp\left(\int_{T/4}^{T/2} dt_1 \mathbf{A}(t_1)\right) \tag{5.37}$$

$$g^{-1}(\tfrac{T}{4})g(\tfrac{3T}{4}) = \prod_{T/2}^{3T/4} \exp(\mathbf{A}(t)dt) \approx \exp\left(\int_{T/2}^{3T/4} dt_1 \mathbf{A}(t_1)\right) \tag{5.38}$$

$$g^{-1}(\tfrac{3T}{4})g(T) = \prod_{3T/4}^{T} \exp(\mathbf{A}(t)dt) \approx \exp\left(\int_{3T/4}^{T} dt_1 \mathbf{A}(t_1)\right) \tag{5.39}$$

But assuming the body velocity is commutative over every quarter of the gait will introduce errors in our analysis, and in this section we quantify these errors. The errors will depend on the connection vector field and the shape of the gait.

To obtain an estimate of the errors introduced by assuming local commutativity, we start with the following assumptions:

1. A linearly varying connection vector field over the shape space.
2. We restrict our analysis to families of gait that trace out squares on the shape space.

We then subdivide each quarter of the gait further and show that the contribution of each quarter calculated using the BCH formula asymptotically approaches, upto $O(l^3)$,

$$g^{-1}(\tfrac{3T}{4})g(T) = \exp\Big(-\beta + \frac{\Delta}{2} - \frac{1}{3}[\Delta_1, \beta]\Big) \tag{5.40}$$

$$g^{-1}(\tfrac{T}{2})g(\tfrac{3T}{4}) = \exp\Big(-\alpha - \frac{\delta}{2} - \frac{1}{3}[\delta_1, \alpha]\Big) \tag{5.41}$$

$$g^{-1}(\tfrac{T}{4})g(\tfrac{T}{2}) = \exp\Big(\beta + \frac{\Delta}{2} + \frac{1}{3}[\Delta_1, \beta]\Big) \tag{5.42}$$

$$g^{-1}(0)g(\tfrac{T}{4}) = \exp\Big(\alpha + \frac{\delta}{2} + \frac{1}{3}[\delta_1, \alpha]\Big) \tag{5.43}$$

where $2l$ is the side length of the square gait centered at the origin of the shape space and,

$$\alpha = 2l\mathbf{A}_1|_{(0,0)} \tag{5.44}$$

$$\beta = 2l\mathbf{A}_2|_{(0,0)} \tag{5.45}$$

$$\delta = 4l^2 \frac{\partial \mathbf{A}_1}{\partial \alpha_2}\bigg|_{(0,0)} \tag{5.46}$$

$$\Delta = 4l^2 \frac{\partial \mathbf{A}_2}{\partial \alpha_1}\bigg|_{(0,0)} \tag{5.47}$$

$$\delta_1 = -l^2 \frac{\partial A_1}{\partial \alpha_1}\bigg|_{(0,0)} \tag{5.48}$$

$$\Delta_1 = -l^2 \frac{\partial A_2}{\partial \alpha_2}\bigg|_{(0,0)} \tag{5.49}$$

$$\tag{5.50}$$

Hence the errors introduced by the local commutativity assumption are of $O(l^3)$ or higher, and more importantly, combining the contribution from each quarter using the BCH formula leads to errors of $O(l^4)$ or higher in the total Lie bracket formula. Explicit details of these calculations are shown in appendix D. Thus errors from assumption II discussed in the next section, which are of $O(l^3)$, would become significant before errors from local commutativity assumption.

5.2.2 Accuracy of Assumption II

In this subsection, we look at the accuracy of the approximation (5.34). Given a, b, c and d such that (5.20)- (5.23) holds, assumption (5.34) states that

$$\exp(a)\exp(b)\exp(c)\exp(d) \approx \exp(-\delta + \Delta + [\alpha, \beta]) \tag{5.51}$$

The expression on the right hand side is obtained by limiting ourselves to terms of $O(l^2)$, where l is the perimeter of the gait. In this section we look at terms of $O(l^3)$.

The order $O(l^3)$ term of the expansion in (5.34) is $\frac{1}{2}[(\alpha+\beta), (-\delta+\Delta+[\alpha,\beta])]$. Thus the leading order error term is

$$E = \frac{1}{2}[(\alpha+\beta), (-\delta+\Delta+[\alpha,\beta])] \tag{5.52}$$

$$= \frac{1}{2}\left[(\alpha+\beta), \iint D\mathbf{A}\right]. \tag{5.53}$$

From (5.13)-(5.16) and (5.20)-(5.23), we get that,

$$\alpha + \beta = \frac{(a+b)-(c+d)}{2} = \frac{1}{2}\left(\int_0^{\frac{T}{2}} dt_1 \mathbf{A}(t_1) - \int_{\frac{T}{2}}^{T} dt_1 \mathbf{A}(t_1) \right). \tag{5.54}$$

Therefore, for gaits belonging to the same image family, unlike the $O(l^2)$ terms, the leading order error terms, of $O(l^3)$, depend on the starting point of the gait.

5.3 Minimizing Leading Order Error

We can use the relation between leading order error and starting point of the gait encoded in (5.53) to identify starting points for an image family of gaits which minimize the leading order error. Starting a gait from a point such that the tangent to the gait at that point coincides with the general direction of the $\mathbf{A}^\theta$ vector field flow in the region enclosed by the gait, reduces the leading order error to zero when the system has a constant rotational vector field. The accuracy of the heuristic goes down with increasing non-linearity of the $\mathbf{A}^\theta$ connection vector field. Justification for this heuristic is presented in appendix E.

In Fig. 5.1, we demonstrate the accuracy of the starting point selection heuristic for the differential drive car (which has a constant rotational vector field) and kinematic snake (which has a non-constant rotational vector field). In Fig. 5.1(a) and Fig. 5.1(c), we show a square image family of gaits with side length, $L = 1$, overlaid on the rotational connection vector field for the differential drive car in the original and optimal coordinates respectively. Similar rotational connection

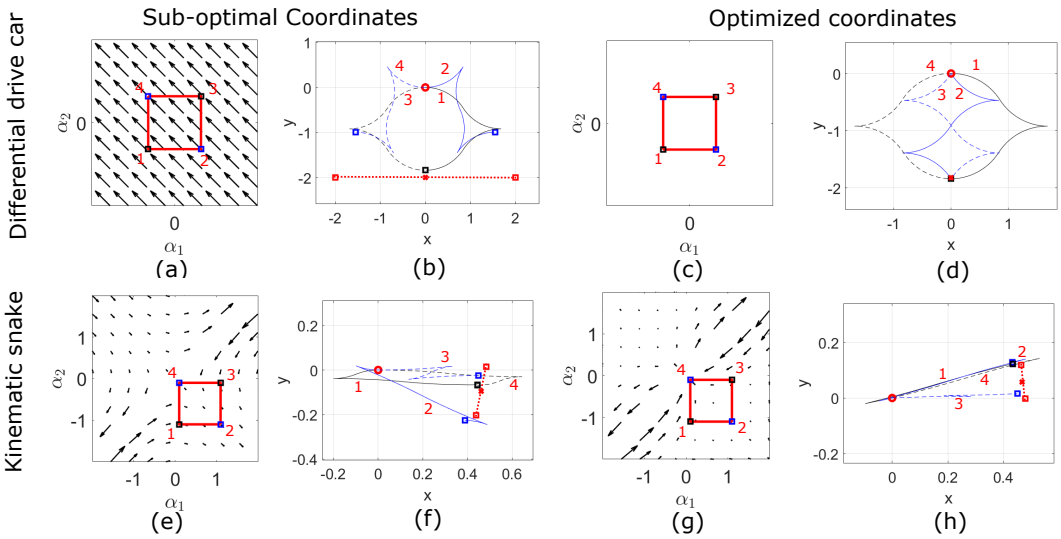

Figure 5.1: A square image family of gaits with side length, $L = 1$, overlaid on the rotational connection vector field for the differential drive car in the original and optimal coordinates are shown in (a) and (c) respectively. Similar rotational connection vector field plots for the kinematic snake are shown in (e) and (g) respectively. The black boxes with numbers 1 and 3 indicate the starting points which minimize the leading order error and the blue boxes with numbers 2 and 4 indicate the starting points which maximize the leading order error. The displacement produced by executing the square gait from these starting points are shown in the figures adjacent to the connection vector field plots. The red circles in (b),(d),(f) and (g) show where the systems starts out from and the red star shows the displacement predicted by the total Lie bracket formula. The magnitude of leading order error E depends on the starting point of the gait. The red dashed lines represent the range of leading order error E. Starting the gait from locations denoted by black boxes produce the least amount of error, while starting from blue boxes maximizes the error as expected. The range of leading order error E is also smaller in the optimal coordinates than in the original coordinate used to describe the systems.

vector field plots for the kinematic snake are shown in Fig. 5.1(e) and Fig. 5.1(g) respectively. In Fig. 5.1(a),(c),(e) and (g), the black boxes with numbers 1 and 3 indicate the starting points which minimize the leading order error and the blue boxes with numbers 2 and 4 indicate the starting points which maximize the leading order error. The displacements produced by executing the square gait from these starting points are shown in the figures adjacent to the connection vector field plots. The red circles in Fig. 5.1(b),(d),(f) and (g) show where the systems starts out from and the red star shows the displacement predicted by the total Lie bracket formula. The magnitude of leading order error E depends on the starting point of the gait. The red dashed lines represent the range of leading order error E.

Starting the gait from locations denoted by black boxes produce the least amount of error, while starting from blue boxes maximizes the error as expected in the differential drive car. In the case of the kinematic snake, our predictions about starting points that produce extremas of leading order error E are accurate despite non-constant $\mathbf{A}^{\theta}$ connection vector fields in the original and optimal coordinates because of symmetries in the rotational vector fields. The range of leading order error E is also smaller in the optimal coordinates than in the original coordinate used to describe the systems.

5.4 Conclusion

Identifying optimal gaits for locomotion is an important topic of research in the robotics community. In [59], we presented a geometric framework for identifying

optimal gaits for kinematic locomoting system. That work built on several results from the geometric mechanics community aimed at relating the displacement produced by a gait to physical parameters of the gait using a truncated BCH expansion. The accuracy of the truncated BCH expansion depended on the choice of body coordinates. In [31], we presented a way of choosing optimal coordinates.

In this paper, our goal was to relate the magnitude of the higher order BCH terms associated with a gait to physical parameters of the gait and the magnitude of the connection vector fields, and show that in the optimal coordinates these decay faster than in the unoptimized coordinates. Along with identifying optimal gaits using the framework, we can now provide an estimate of uncertainty of resulting displacement because of ignoring higher order BCH expansion terms.

Chapter 6: Data-driven Connection in Granular Media

We have focused on swimmers in drag dominated-environments so far in this thesis. The dynamics governing the operation of these systems can be analytically obtained using resistive force theory (RFT). The optimization framework presented in this thesis is applicable to a larger class of kinematic systems. In this chapter, we demonstrate the process of obtaining a data-driven connection for a robot operating in granular media. This will lay the groundwork for future exploration into identifying optimal gaits in granular media using the framework presented in this thesis.

Studying motion of animals and robots that maneuver through sand, rubble and debris is often analytically intractable and computationally very expensive. This makes obtaining an analytical model or a high-fidelity finite element model difficult.

Previous research has shown that for sand swimming of robots in dry granular media, arguably the simplest flowing terrestrial material, the granular media can be modelled as a "frictional fluid" in which forces are dominated by Coulomb friction, making them insensitive to rate, and in which inertial effects are small. Hence we would expect the connection vector fields generated by our approach for the SEA snake robot used for experiments in this chapter to be qualitatively similar to the connection vector field of a swimming robot with an identical morphology in a

drag-dominated fluid.

6.1 Previous Work

A simple Purcell swimmer in granular media was studied in [27]. An empirical swimming model of the robot was generated by treating the media as a "frictional fluid." The predictions of displacements produced by the robot were shown to be excellent at small amplitudes and good at large amplitudes.

A data-driven geometric model of a snake robot was presented in [16] by repeating the following process over a uniform sampling of the shape space.

1. Sampling a configuration
2. Commanding a small shape velocity
3. Measuring the resultant body velocity
4. Fitting the local connection

When experimentally determining the local connection, the surface is flattened after each "twitch". This leads to a contact pattern of the robot and media that is different from what the robot actually experiences when executing a gait, resulting in a discrepancy in the magnitude of predicted and observed displacements.

6.2 Data-driven Connection Generation

6.2.1 Shape and position data from encoders and markers

The data for generating the connection for the SEA snake is obtained by letting the robot do multiple trials of different gaits (circular, spiral, rotating ellipsoidal). In Fig. 6.1, we overlay shape space trace from all the trials used for obtaining the data driven connection vector field. We can see it provides a good covering of the shape space.

The next step to obtaining the data driven connection is deciphering the shape velocity and position velocity from marker and encoder data at each data point. Encoder data from the SEA robot is fairly robust and obtaining shape and shape velocity data from encoder data is fairly straightforward. During data collection, some markers were often occluded by the granular medium and the encoder data is generally more accurate than the marker data. We use a combination of marker and joint angle data to obtain the body frame, which is defined to be located at the center of the fifth link and oriented along the fifth link. At each step we use the encoder data to correct the marker data so that the marker data respects the observed joint angles.

The 7 markers for data collection are placed approximately at the centers of links 2,3,4,5,6,7 and 8. Let (x_i, y_i) denote the location of the markers on link i. If we assume the the center of the 5th link is at $(\hat{x}_5, \hat{y}_5)$ and the link is oriented at angle θ, we can reconstruct the locations of the centers of other links through kinematics and joint angle data. Let's call these locations $(\hat{x}_i, \hat{y}_i)$. We find the

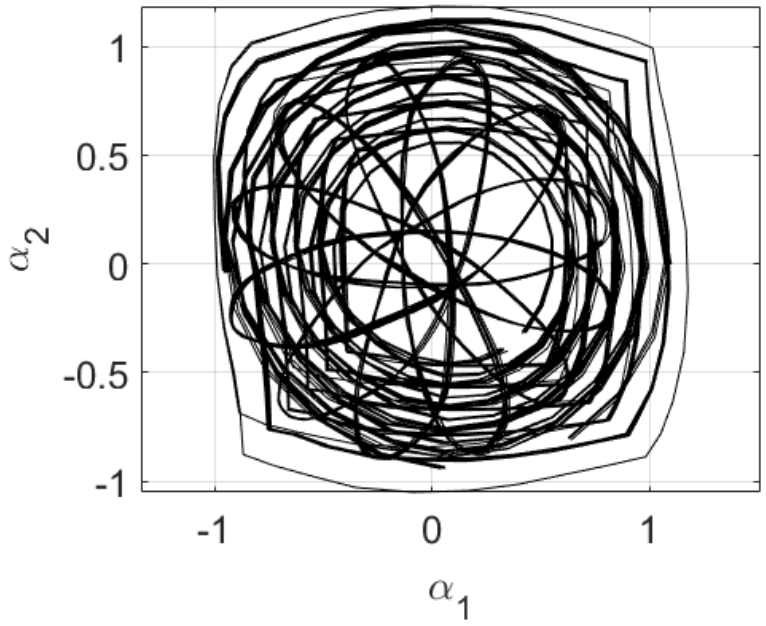

Figure 6.1: Motion of the SEA snake from all datasets imposed on a single figure. These motions were used to experimentally determine the local connection **A** *for the SEA snake.*

position and orientation of the body frame by running the following optimization at each time step of each trial:

$$\underset{\hat{x}_5,\hat{y}_5,\theta}{\operatorname{argmin}} \sum_{i=2}^{8} (x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2 \tag{6.1}$$

This process is equivalent to correcting the observed marker data at each step so that it respects the joint angle data from the encoders and then obtaining the location and orientation of the middle link from the corrected marker data.

6.2.2 Connection from shape and position data

The process of finding the data-driven connection vector field from the shape velocity and position data is identical to the method presented in [62]. As explained in Chapter 2, the geometric model assumes a linear relationship between the system's shape and changes in its position,

$$\mathring{g} = \mathbf{A}(r)\dot{r}. \tag{6.2}$$

The local connection $\mathbf{A}$ is a function of purely the shape of the robot and the goal of our analysis is to quantify this relationship. We do this by

1. Overlaying a uniform grid on top of our shape space
2. Estimating the value of the local connection at each grid point using data points (combination of shape velocity and the resulting position velocity) near the grid point [62].

6.3 Results

The connection vector fields obtained through data-driven analysis are shown in Fig. 6.2. The y and θ-directional vector fields agree nicely with what we expect from theory. The x- directional vector field shows qualitatively similar properties to the vector field predicted by theory but the agreement is not as good as that seen in y and θ directions. In Fig. 6.3,6.4 and 6.5, we compare the body velocities predicted by our data-driven vector fields (black-dashed lines) with the observed

body velocities (colored lines) in the x, y and θ direction respectively.

6.4 Future Directions

6.4.1 Predicting displacements for systems with higher dimensional shape spaces

The connection vector field for the viscous swimmer was generated using Resistive Force Theory (RFT) with a 2:1 ratio between normal and tangential drag coefficients. We know RFT describes granular locomotion well from previous work but the normal to tangential drag coefficient depends on the material of the robot and the substrate. If we can identify this ratio from the data driven vector fields we have generated, we can use it to predict the displacements produced by the SEA snake when we use more than two shape modes to describe its motion.

The brute force way of finding the normal to tangential drag ratio governing the SEA snake would be to generate theoretical vector fields at a range of these ratios and identify the vector field that matches best with the data-driven vector field. Combining the data-driven physics framework with the soap bubble optimizer, we would be able to predict optimal gaits when complex and more than two shape modes are employed from data obtained while executing gaits described by simple shape modes.

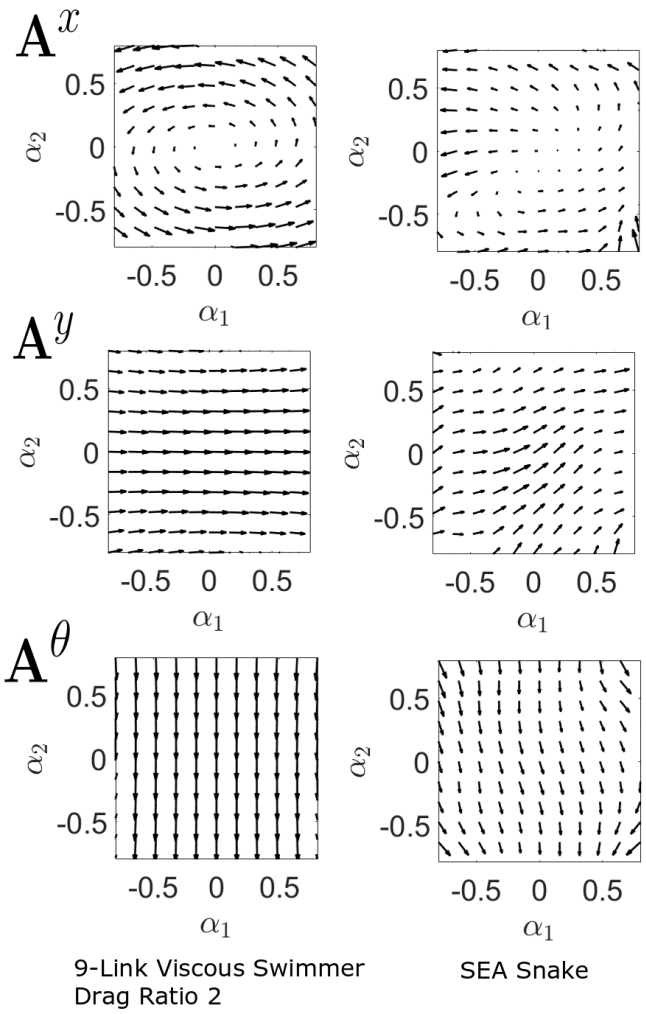

Figure 6.2: Comparison of the connection vector fields obtained through RFT for the 9 link viscous swimmer with the connection vector fields obtained through data driven analysis for the SEA snake. The connection vector fields show many qualitative similarities.

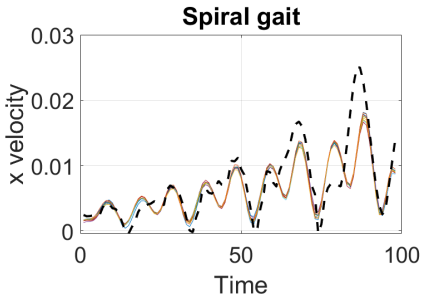

Figure 6.3: A comparison of the predicted x-directional body velocity by our data-driven vectorfield (black dotted line) and the observed x-directional body velocity when executing spiral gaits.

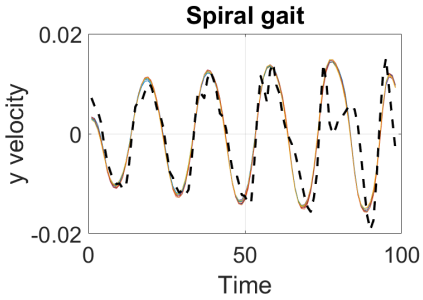

Figure 6.4: A comparison of the predicted y-directional body velocity by our data-driven vectorfield (black dotted line) and the observed y-directional body velocity when executing spiral gaits.

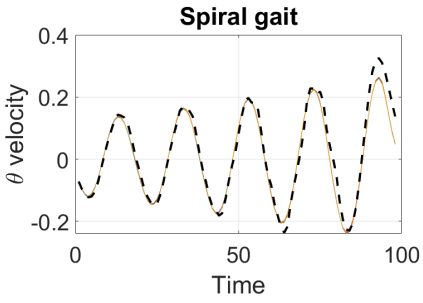

Figure 6.5: A comparison of the predicted rotational body velocity by our data-driven vectorfield (black dotted line) and the observed rotational body velocity when executing spiral gaits.

6.4.2 Inclines

Another interesting line of inquiry would be generating the data driven vector fields when the SEA snake is operating on an inclined surface. Along with the normal to tangential drag ratio, the drag would also depend on whether the snake is moving up or down the incline. An initial study could involve identifying the connection vector fields when the snake is moving purely downhill and when the snake is moving purely uphill. It would be interesting to see if the difference in dynamics between moving uphill and downhill can be captured just by a change in the normal to tangential drag ratio or if we need a physics model fundamentally different from RFT to capture the dynamics of the snake moving on an incline.

Chapter 7: Conclusion

In this thesis, we consolidate upon and extend prior geometric insights about locomotion into a set of geometric principles that govern the shape of optimally-efficient gaits for drag dominated systems. We formally encode these principles in a set of geometric expressions that together make up the gradient of the gait efficiency with respect to variations in the gait trajectory. We use this gradient in a gradient-descent solver to find optimal gaits, but more fundamentally, *the gradient geometrically describes the dynamics underlying any other gait optimization algorithm applied to the system.*

For systems with two shape variables, the dynamics of this solver for a fully actuated swimmer are analogous to those of a soap bubble, with the Lie bracket providing an "inflating pressure" to the trajectory and the Riemannian metric on the shape space contributing "surface tension" that halts growth of the cycle in the face of diminishing returns, and a "concentration gradient" that provides a power-optimal pacing along the gait. Together, these elements drive the gait cycle to a "comfortable stride" that converts shape change effort into net displacement with optimal efficiency.

By extending the gradient calculation to systems with three shape variables we see that dynamics of our solver generalize to maximizing flux through an oriented surface. For systems with three shape variables, the constraint curvature functions

can be visualized as vector flux because each surface element has a unique normal.

In systems with more than three shape variables, each surface element no longer has a unique normal direction associated with it, and we thus drop the "vector flux" analogy for these higher-dimensional systems, but can preserve the idea of two-form flux passing through the surface elements that make up the "interior" of the gait.

We demonstrated this variational principle in operation on a number of test systems in viscous-dominated environments, including Purcell's three-link swimmer (a standard minimal template for locomotion modeling) and a serpenoid swimmer (a model widely used in studies of animals and snake robots). In the lower-dimensional cases, the optimal gaits found by our approach match those previously found by exhaustive optimizations of the gait cycles [66], and in the higher-dimensional cases, the optimizer allowed us to efficiently explore a space of candidate swimming morphologies.

We also presented how the framework can be used to simultaneously optimize design and control variables for locomoting systems. We demonstrated this by finding the optimal link length ratio for the Purcell swimmer.

In the context of related works, our framework can be viewed as a "macroscopic" extension of the Lie-bracket control schemes in [50, 52]. This extension is significant because it uses the geometry of the systems' dynamics to identify the *amplitudes* of their most efficient gait cycles (in addition to the useful sets of phase-couplings identified in previous work). For our example systems swimming at low Reynolds number, the two-dimensional "soap-bubble" analysis can be seen

as an instantiation of the boundary-value problem suggested in [3], for which we have analytically constructed a gradient from the curvature of the constraints, and the higher-dimensional portions of a work to be a generalization of this principle. These extensions and generalizations of previous work in the literature have both been enabled by our calculation of a minimum-perturbation body frame [34], which significantly increases the accuracy of Lie-bracket approximations to large motions.

We then extend this framework to study the geometric structure of optimal gaits for viscous swimmers with passive elastic joints by combining the constraint-curvature analysis in [61] with frequency-response models for the steady state motion of driven oscillators. We use this structure to identify both speed-maximizing and efficiency-maximizing gaits. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

As discussed in §4.2.2, for the fully actuated swimmers, the maximum forward speed achievable is only restricted by the maximum power we are able to supply the joints, but for the swimmers with the passive elastic joint, even with more powerful actuators, there is a theoretical maximum forward speed the system can achieve dictated by the stiffness of the passive joint.

The important factor that makes the performance of the fully actuated swimmers superior to that of the swimmers with the passive elastic joint in terms of energy efficiency is the fact that not only can the fully actuated swimmers execute a much larger set of gaits, they can execute any gait the passive swimmer can

execute at a pacing that is just as good or better than the pacing dictated by the dynamics of the passive joint.

This raises the question of what benefits, e.g., simplicity of construction, does having a passive elastic member give to biological organisms that locomote in a low Reynolds number fluid? Most biological organisms have tails that resemble an elastic filament. The propulsive and flexive dynamics of such filaments have been well studied [13, 25, 69, 70]. Artificial microscopic swimmers with elastic filaments have been proposed based on this body of work [19, 20]. An interesting line of future work would involve investigating the tradeoff between elastic element inefficiencies and structural complexity of being fully actuated.

The work presented in chapter 4 is the first step towards expanding the applicability of the geometric framework presented in [59–61] to systems where underactuated shape parameters play a role in the dynamics of the system. In the case of the passive Purcell swimmer, assuming the torque required to affect a desired shape change did not depend on the current shape of the swimmer did not introduce significant errors in the predicition of the limit cycle corresponding to inputs to the actuated joint as shown in Fig. 4.1. This might not always be the case in all the swimmers we consider. A future line of research would be to improve our frequency domain analysis by using non-linear perturbation theory to obtain more accurate predictions of limit cycles.

The gradient calculation process in Chapters 3 and 4 depends on the total Lie bracket formula obtained by restricting the BCH expansion to the first few terms. It is crucial that the truncation be a good approximation of the full BCH

expansion. In chapter 5 we explain the factors affecting the accuracy of the total Lie bracket formula. We show that the major factors affecting the accuracy are the magnitude of the connection vector field and the size of the gait. Since, the optimal coordinate choice minimizes the magnitude of the connection vector field, it improves the accuracy of the truncated Magnus series and hence the accuracy of the total Lie bracket formula. We relate the accuracy of these assumptions to the system dynamics and the physical parameters of the gait, present a heuristic for selecting a starting point that minimizes the leading order error associated with the total Lie bracket formula.

A key goal of our geometric study of locomotion is to make this approach feasible for robots whose dynamics are difficult to model from first principles, as in the granular systems of [16, 27]. As they do for the ideal system models used as examples in this paper, the dimensionality extensions and explicit optimality criteria discussed here will broaden the classes of systems that we can consider in the empirical-geometric work. Additionally, the geometric structure we discuss here forms the foundation of our geometric-Floquet approach to online data-driven locomotion analysis and optimization on very-high dimensional systems [10].

Bibliography

[1] F Alouges, A DeSimone, and A Lefebvre. Optimal strokes for axisymmetric microswimmers. *The European Physical Journal E*, 28(3):279–284, 2009.

[2] François Alouges, Antonio DeSimone, and Aline Lefebvre. Optimal strokes for low reynolds number swimmers: an example. *Journal of Nonlinear Science*, 18(3):277–302, 2008.

[3] J. E. Avron, O. Gat, and O. Kenneth. Optimal swimming at low reynolds numbers. *Physical Review Letters*, 93(18):186001–1–186001–r, 2004.

[4] Joseph E. Avron and Oren Raz. A geometric theory of swimming: Purcell's swimmer and its symmetrized cousin. *New Journal of Physics*, 9(437), 2008.

[5] C Bär. *Carnot-Caratheodory-Metriken*. PhD thesis, Diplomarbeit, Bonn, 1988.

[6] C Bär. Geodesics for carnot-caratheodory metrics. *Preprint*, 1989.

[7] L.E. Becker, S. A. Koehler, and H. A. Stone. On self-propulsion of micro-machines at low Reynolds number: Purcell's three-link swimmer. *Journal of Fluid Mechanics*, 490:15–35, 2003.

[8] Piernicola Bettiol, Bernard Bonnard, Laetitia Giraldi, Pierre Martinon, and Jérémy Rouot. The Purcell Three-link swimmer: some geometric and numerical aspects related to periodic optimal controls. working paper or preprint, October 2015.

[9] Piernicola Bettiol, Bernard Bonnard, and Jérémy Rouot. Optimal strokes at low reynolds number: a geometric and numerical study of copepod and purcell swimmers. *SIAM Journal on Control and Optimization*, 56(3):1794–1822, 2018.

[10] Brian Bittner, Ross L Hatton, and Shai Revzen. Geometrically optimal gaits: a data-driven approach. *Nonlinear Dynamics*, pages 1–16, 2018.

[11] Anthony M. Bloch et al. *Nonholonomic Mechanics and Control*. Springer, 2003.

[12] Callie Branyan, Chloë Fleming, Jacquelin Remaley, Ammar Kothari, Kagan Tumer, Ross L Hatton, and Yiğit Mengüç. Soft snake robots: Mechanical design and geometric gait implementation. In *Robotics and Biomimetics (ROBIO), 2017 IEEE International Conference on*, pages 282–289. IEEE, 2017.

[13] Christopher Brennen and Howard Winet. Fluid mechanics of propulsion by cilia and flagella. *Annual Review of Fluid Mechanics*, 9(1):339–398, 1977.

[14] A. E. Bryson and Y.-C. Ho. *Applied Optimal Control, Optimization, Estimation, and Control*. John Wiley & Sons, New York-London-Sydney-Toronto, 1975.

[15] Lisa J. Burton, Ross L. Hatton, Howie Choset, and Anette E. Hosoi. Two-link swimming using buoyant orientation. *Physics of Fluids*, 22:091703, September 2010.

[16] Jin Dai, Hossein Faraji, Chaohui Gong, Ross L. Hatton, Daniel I. Goldman, and Howie Choset. Geometric swimming on a granular surface. In *Proceedings of the Robotics: Science and Systems Conference*, pages 1–8, Ann Arbor, Michigan, July 2016.

[17] Tony Dear, Ross L. Hatton, and Howie Choset. Nonlinear dimensionality reduction for kinematic cartography with an application toward robotic locomotion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Chicago, IL, September 2014.

[18] Antonio DeSimone, Luca Heltai, François Alouges, and Aline Lefebvre-Lepot. Computing optimal strokes for low Reynolds number swimmers. *Natural Locomotion in Fluids and on Surfaces*, pages 177–184, 2012.

[19] Rémi Dreyfus, Jean Baudry, Marcus L Roper, Marc Fermigier, Howard A Stone, and Jérôme Bibette. Microscopic artificial swimmers. *Nature*, 437(7060):862, 2005.

[20] Jon Edd, Sébastien Payen, Boris Rubinsky, Marshall L Stoller, and Metin Sitti. Biomimetic propulsion for a swimming surgical micro-robot. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, volume 3, pages 2583–2588. IEEE, 2003.

[21] Michael J Enos. *Dynamics and control of mechanical systems: the falling cat and related problems*, volume 1. American Mathematical Soc., 1993.

[22] Harley Flanders. Differentiation under the integral sign. *The American Mathematical Monthly*, 80(6):615–627, 6 1973.

[23] L. Giraldi, P. Martinon, and M. Zoppello. Controllability and optimal strokes for n-link microswimmer. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 3870–3875, Dec 2013.

[24] Nicola Giuliani, Luca Heltai, and Antonio DeSimone. Predicting and optimizing microswimmer performance from the hydrodynamics of its components: The relevance of interactions. *Soft Robotics*, 0(0):null, 0. PMID: 29762082.

[25] Raymond E Goldstein and Stephen A Langer. Nonlinear dynamics of stiff polymers. *Physical review letters*, 75(6):1094, 1995.

[26] Chaohui Gong, Daniel I. Goldman, and Howie Choset. Simplifying gait design via shape basis optimization. In *Proceedings of Robotics: Science and Systems*, AnnArbor, Michigan, June 2016.

[27] R. L. Hatton, H. Choset, Y. Ding, and D. I. Goldman. Geometric visualization of self-propulsion in a complex medium. *Physical Review Letters*, 110:078101, February 2013.

[28] Ross L. Hatton, Lisa J. Burton, Anette E. Hosoi, and Howie Choset. Geometric maneuverability, with applications to low Reynolds number swimming. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, CA USA, September 2011.

[29] Ross L. Hatton and Howie Choset. Approximating displacement with the body velocity integral. In *Proceedings of Robotics: Science and Systems V*, Seattle, WA USA, June 2009.

[30] Ross L. Hatton and Howie Choset. Optimizing coordinate choice for locomoting systems. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4493–4498, Anchorage, AK USA, May 2010.

[31] Ross L. Hatton and Howie Choset. Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice. *International Journal of Robotics Research*, 30(8):988–1014, July 2011.

[32] Ross L. Hatton and Howie Choset. Kinematic cartography for locomotion at low Reynolds numbers. In *Proceedings of Robotics: Science and Systems VII*, Los Angeles, CA USA, June 2011.

[33] Ross L. Hatton and Howie Choset. Geometric swimming at low and high Reynolds numbers. *IEEE Transactions on Robotics*, 29(3):615–624, June 2013.

[34] Ross L. Hatton and Howie Choset. Nonconservativity and noncommutativity in locomotion. *European Physical Journal Special Topics: Dynamics of Animal Systems*, 224(17–18):3141–3174, 2015.

[35] Ross L. Hatton and Howie Choset. Kinematic cartography and the efficiency of viscous swimming. *IEEE Transactions on Robotics*, 33(3):523–535, June 2017.

[36] Ross L Hatton, Tony Dear, and Howie Choset. Kinematic cartography and the efficiency of viscous swimming. *IEEE Transactions on Robotics*, 2017.

[37] S. Hirose. *Biologically Inspired Robots (Snake-like Locomotor and Manipulator)*. Oxford University Press, 1993.

[38] Ikhee Jo, Yangyang Huang, Walter Zimmermann, and Eva Kanso. Passive swimming in viscous oscillatory flows. *Physical Review E*, 94(6):063116, 2016.

[39] Eva Kanso. Swimming due to transverse shape deformations. *Journal of Fluid Mechanics*, 631:127–148, 2009.

[40] Scott D. Kelly Kelly and Richard M Murray. Geometric phases and robotic locomotion. *J. Robotic Systems*, 12(6):417–431, Jan 1995.

[41] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of Differential Geometry*, volume 1. Wiley Interscience, 1963.

[42] Deepak Krishnamurthy, Georgios Katsikis, Arjun Bhargava, and Manu Prakash. Schistosoma mansoni cercariae swim efficiently by exploiting an elastohydrodynamic coupling. *Nature Physics*, 13(3):266, 2017.

[43] Zexiang Li and John F Canny. *Nonholonomic motion planning*, volume 192. Springer Science & Business Media, 2012.

[44] James Lighthill. Flagellar hydrodynamics. *SIAM Review*, 18(2):161–230, 1976.

[45] Wilhelm Magnus. On the exponential solution of differential equations for a linear operator. *Communications on Pure and Applied Mathematics*, VII:649–673, 1954.

[46] Juan B. Melli, Clarence W. Rowley, and Dzhelil S. Rufat. Motion planning for an articulated body in a perfect planar fluid. *SIAM Journal of Applied Dynamical Systems*, 5(4):650–669, November 2006.

[47] Richard Montgomery. Isoholonomic problems and some applications. *Communications in Mathematical Physics*, 128(3):565–592, 1990.

[48] Richard Montgomery. *A tour of subriemannian geometries, their geodesics and applications*. Number 91. American Mathematical Soc., 2006.

[49] Alessandro Montino and Antonio DeSimone. Three-sphere low-reynolds-number swimmer with a passive elastic arm. *The European Physical Journal E*, 38(5):42, 2015.

[50] Kristi A. Morgansen, Benjamin I. Triplett, and Daniel J. Klein. Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles. *IEEE Transactions on Robotics*, 23(6):1184–1199, Jan 2007.

[51] Ranjan Mukherjee and David P. Anderson. A surface integral approach to the motion planning of nonholonomic systems. In *American Control Conference, 1993*, pages 1816 –1823, 1993.

[52] Richard M. Murray and Shankar S. Sastry. Nonholonomic motion planning: Steering using sinusoids. *IEEE Transactions on Automatic Control*, 38(5):700–716, Jan 1993.

[53] James P. Ostrowski and Joel Burdick. The mechanics and control of undulatory locomotion. *International Journal of Robotics Research*, 17(7):683–701, July 1998.

[54] James P. Ostrowski, Jaydev P. Desai, and Vijay Kumar. Optimal gait selection for nonholonomic locomotion systems. *International Journal of Robotics Research*, 19(3):225–237, 2000.

[55] J.P. Ostrowski and J.W. Burdick. The geometric mechanics of undulatory robotic locomotion. *International Journal of Robotics Research*, 17(7):683–702, 1998.

[56] Emiliya Passov and Yizhar Or. Dynamics of purcell's three-link microswimmer with a passive elastic tail. *The European Physical Journal E*, 35(8):1–9, 2012.

[57] E. M. Purcell. Life at low Reynolds numbers. *American Journal of Physics*, 45(1):3–11, January 1977.

[58] James E. Radford and Joel W. Burdick. Local motion planning for nonholonomic control systems evolving on principal bundles. In *Proceedings of the International Symposium on Mathematical Theory of Networks and Systems*, Padova, Italy, 1998.

[59] Suresh Ramasamy and Ross L Hatton. Soap-bubble optimization of gaits. In *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pages 1056–1062. IEEE, 2016.

[60] Suresh Ramasamy and Ross L Hatton. Geometric gait optimization beyond two dimensions. In *American Control Conference (ACC), 2017*, pages 642–648. IEEE, 2017.

[61] Suresh Ramasamy and Ross L Hatton. The geometry of optimal gaits for drag-dominated kinematic systems. *IEEE Transactions on Robotics*, 35(4):1014–1033, 2019.

[62] Jacquelin L Remaley. Improving understanding of geometric swimmer locomotion. 2018.

[63] Guy Rodnay and Elon Rimon. Isometric visualization of configuration spaces of two-degrees-of-freedom mechanisms. *Mechanism and machine theory*, 36(4):523–545, 2001.

[64] Elie A. Shammas, Howie Choset, and Alfred A. Rizzi. Geometric motion planning analysis for two classes of underactuated mechanical systems. *Int. J. of Robotics Research*, 26(10):1043–1073, 2007.

[65] Alfred Shapere and Frank Wilczek. Geometry of self-propulsion at low Reynolds number. *Journal of Fluid Mechanics*, 198:557–585, 1989.

[66] Daniel Tam and Anette E. Hosoi. Optimal stroke patterns for Purcell's three-link swimmer. *Phys. Review Letters*, 98(6):068105, 2007.

[67] J.J. van Wijk and A.M. Cohen. Visualization of Seifert surfaces. *Visualization and Computer Graphics, IEEE Transactions on*, 12(4):485–496, 2006.

[68] Gregory C. Walsh and Shankar Sastry. On reorienting linked rigid bodies using internal motions. *Robotics and Automation, IEEE Transactions on*, 11(1):139–146, January 1995.

[69] Chris H Wiggins and Raymond E Goldstein. Flexive and propulsive dynamics of elastica at low reynolds number. *Physical Review Letters*, 80(17):3879, 1998.

[70] Tony S Yu, Eric Lauga, and AE Hosoi. Experimental investigations of elastic tail propulsion at low reynolds number. *Physics of Fluids*, 18(9):091701, 2006.

[71] Lijun Zhu, Zhiyong Chen, and Tetsuya Iwasaki. Oscillation, orientation, and locomotion of underactuated multilink mechanical systems. *IEEE Transactions on Control Systems Technology*, 21(5):1537–1548, 2013.

APPENDICES

Appendix A: Relation to sub-Riemannian Geometry

In the geometric mecahnics community, optimal paths between two points in the configuration spaces of swimmers in drag dominated environments have been identified as being the subriemannian geodesics joining these two points [48]. In this section, we review what subriemannian geodesics are, discuss previous research aimed at finding these subriemannian geodesics and present why the variational framework presented in this work simplifies the process of finding a subriemannian geodesic.

For swimmers in a drag dominated environment, the constraints (which are completely encoded in the local connection $\mathbf{A}$) determine a vector subbundle of the tangent bundle of Q, $\mathcal{H} \subset TQ$ called the horizontal distribution. This horizontal distribution represents the set of all velocities $\dot{q}$ the swimmer can achieve. Any path that is tangent to $\mathcal{H}$ is called a horizontal path. Any absolutely continuous horizontal path that minimizes the distance (according to the power metric $\mathcal{M}$) between any points on it is called a subriemannian geodesic [48].

Many methods to find these geodesics are based on the theorem proved in [5, 6], which states that every subriemannian geodesic is the cotangent projection to Q of a solution on T^*Q to the Hamiltonian differential equations for the Hamiltonian H, which is uniquely determined by the power metric $\mathcal{M}$ and the horizontal distribution $\mathcal{H}$. In general the Hamiltonian differential equations are a set of $2n$

coupled first order differential equations, where n is the dimension of Q. This makes calculating the subriemmanian geodesic a numerically expensive problem. A shooting based method to find subriemannian geodesics is presented in [2].

Gaits are closed loops in the shape space R, which can be identified as the quotient space of Q under the group action G, Q/G. With each gait $\phi : [0, 1] \to R$ and a base point $\phi(0)$, we can associate a measure of the movement along the fiber produced by executing the gait, called the representative holonomy of ϕ with respect to $\phi(0)$. The problem of finding the shortest loop with a given holonomy is called the isoholonomic problem and has been studied in detail in [21, 43, 48]. The version of this problem most relevant to us is finding a shortest loop with a given holonomy class independent of the base point in R as studied in [47]. This shortest loop is a subriemannian geodesic, and [2] presents a shooting based method to find such a geodesic. This method was also used in [1] to identify optimal gaits for low-Re swimmers. That paper also presented analytical solutions for small shape changes that optimize efficiency. Similar schemes were used in [9] to identify optimal gaits for copepod swimmer and the purcell swimmer. The problem addressed in this work can be considered as an extended isoholonomic problem where amongst all the subriemannian geodesics that are horizontal lifts of closed loops in the shape space and have representative holonomy in a certain desired direction, we want to find the geodesic that maximizes the efficiency η as defined in 2.4. The numerical advantages of our method stem from the facts that:

1. In the optimal coordinates, all gaits within the same image family have approximately the same holonomy. So we only have to identify the image of the

gait, whose horizontal lift yields the most efficient subriemmanian geodesic with holonomy in a certain desired direction.

2. Switching to the optimal coordinates makes the approximation in (5.5) valid for large angle gaits. This approximation enables quicker computation of the change in holonomy with respect to small perturbations of the gait ϕ, because the approximation makes the mapping from the image of the gait to holonomy of the gait history independent (In the optimal coordinates, if we change one part of the gait, the contribution of the other parts towards the holonomy of the gait does not change).

Appendix B: Kinematics of LowRe Swimmers Obtained using RFT

B.1 Kinematics of LowRe Swimmers Obtained Using RFT

For all the example systems in the work, we obtain the local connection $\mathbf{A}$ and metric $\mathcal{M}$ by applying a resistive force model to the system geometry. This appendix describes how to define the geometry for a n-link swimmer and how to apply RFT to obtain $\mathbf{A}$ and $\mathcal{M}$.

B.2 Position of links and joints

Each link λ_i in the chain has link length ℓ_i, and each joint σ_j in the chain has an angle α_j. Using the standard matrix representation of planar translations,

$$(x, y, \theta) = \begin{bmatrix} \cos\theta & -\sin\theta & x \\ \sin\theta & \cos\theta & y \\ 0 & 0 & 1 \end{bmatrix}, \tag{B.1}$$

the transformation associated with moving by half a link length (from the proximal end of a link to its midpoint, or from the midpoint to the distal end) is

$$h_i = \begin{bmatrix} 1 & 0 & \ell_i/2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{B.2}$$

and the transformation associated with each joint (from the distal end of the proximal link to the proximal end of the distal link) is

$$a_j = \begin{bmatrix} \cos\alpha_j & -\sin\alpha_j & 0 \\ \sin\alpha_j & \cos\alpha_j & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{B.3}$$

Taking the midpoint of the first link as being at the origin (with an identity transformation), the relative transformation from this link to each other link is given by the right-propagating product of the intermediate transformations,

$$g_{\lambda_{i=1}} = \mathrm{Id} \tag{B.4}$$

$$g_{\lambda_{i\geq 2}} = \prod_{k=2}^{i} \rightarrow (h_{k-1})(a_k)(h_k), \tag{B.5}$$

where the midpoint-to-midpoint transformations are formed by combining a half-length step along the proximal link, the joint rotation, and a half-length step along the distal link.

The joint locations are calculated similarly, with the stator of the first joint at the end of the first link, and the location of each subsequent joint given by the product of the relative transformations,

$$g_{\sigma_{j=1}} = h_1 \tag{B.6}$$

$$g_{\sigma_{j\geq 2}} = h_1 \prod_{k=2}^{j} (a_{k-1})(h_k)(h_k), \tag{B.7}$$

where the stator-to-stator transformations are formed by combining the joint rotation and two half-steps along the link.

Multiplying the inverse of the joint positions by the link positions gives the positions of the links relative to the joints,

$$h_{\frac{i}{j}} = (g_{\sigma_j}^{-1})(g_{\lambda_i}). \tag{B.8}$$

B.3 Velocities of links and joints

If the joints are held rigid, the body velocity of link i is related to the body velocity of the base link by the adjoint-inverse mapping associated with its position relative to the base link,

$$\mathring{g}_i = Ad_{g_{\lambda_i}}^{-1} \ \mathring{g}_1, \tag{B.9}$$

where the adjoint-inverse mapping encodes the cross-product and rotation required to transfer velocities between frames on a rigid body,

$$Ad_g^{-1} = \overbrace{\begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}}^{\text{rotation}} \overbrace{\begin{bmatrix} 1 & 0 & -y \\ 0 & 1 & x \\ 0 & 0 & 1 \end{bmatrix}}^{\text{cross product}} \tag{B.10}$$

$$= \begin{bmatrix} \cos\theta & \sin\theta & x\sin\theta - y\cos\theta \\ -\sin\theta & \cos\theta & x\cos\theta + y\sin\theta \\ 0 & 0 & 1 \end{bmatrix}. \tag{B.11}$$

Similarly, if the base link is held fixed and a single joint is rotated, the body velocity of each link distal to that joint is related to the joint velocity by the adjoint-inverse mapping associated with the position of the link relative to the joint (and links proximal to the joint do not move),

$$\mathring{g}_{i>j} = Ad^{-1}_{h_{\frac{i}{j}}} \begin{bmatrix} 0 \\ 0 \\ \dot{\alpha}_j \end{bmatrix} \qquad \Leftarrow \text{links distal to joint} \tag{B.12}$$

$$\mathring{g}_{i\leq j} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \qquad \Leftarrow \text{links proximal to joint.} \tag{B.13}$$

If both the base link and the joints are moving, then the body velocity of the links is the sum of the contributions from the base-link and joint motion.

This means that we can combine the adjoint-inverse mappings from (B.9), (B.12), and (B.13) into a set of augmented matrices J_i that serve as Jacobians from the system's generalized velocity (body velocity of the base link and the joint angular velocities) to the body velocities of the links,

$$\overset{\circ}{g}_i = J_i \begin{bmatrix} \overset{\circ}{g}_1 \\ \hline \dot{\alpha}_1 \\ \vdots \\ \dot{\alpha}_{i-1} \\ \hline \dot{\alpha}_i \\ \vdots \\ \dot{\alpha}_m \end{bmatrix} \tag{B.14}$$

where

$$J_i = \left[\begin{array}{c|c|ccc|c} & & & & & \\ Ad^{-1}_{g_{\lambda_i}} & Ad^{-1}_{h_{\frac{i}{1}}} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} & \cdots & Ad^{-1}_{h_{\frac{i}{i-1}}} & \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} & \mathbf{0}^{3\times(n-i)} \\ & & & & & \end{array}\right]. \tag{B.15}$$

B.4 Low Reynolds number resistive force model

With the kinematics of the system described in (B.5) and (B.14), we can calculate the forces on the system. At low Reynolds number, swimming bodies experience linear resistive drag, with geometry-dependent coefficients. In a full model for these forces, the drag coefficients depend on the relative proximity and orientation

of the bodies, but a useful approximation to the dynamics for slender bodies is to take the drag forces on a link as being linearly related to the link's body velocity as

$$F_i^b = -\overbrace{\begin{bmatrix} L & & \\ & kL & \\ & & \frac{kL^3}{12} \end{bmatrix}}^{D} \mathring{g}_i, \tag{B.16}$$

where k is the ratio between longitudinal and lateral drag coefficients and the $\frac{kL^3}{12}$ term is the result of integrating the lateral drag on a spinning object over its length, $\int_{-L/2}^{L/2} (ks)s\, ds$. (In a more general model, the structure of (B.16) would be preserved, but D would become a function of α and would potentially gain off-diagonal terms.)

Forces are mapped along bodies by the *dual adjoint* actions Ad_g^*, which are encoded by the transposes of the adjoint actions. The map from body-frame forces on the ith link, F_i^b, to the corresponding body-frame forces acting on the base link, F_1^b, is specifically given by the dual adjoint-inverse mapping,

$$F_1^b = Ad_{g_{\lambda_i}}^{-*} F_i^b, \tag{B.17}$$

which expands as

$$Ad_g^{-*} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ x\sin\theta - y\cos\theta & x\cos\theta + y\sin\theta & 1 \end{bmatrix} \tag{B.18}$$

and, as before, encodes the rotation and cross product operations. Here, because of the transposition, the dual adjoint action transforms translational forces in the input into rotational moments in the output (rather than transforming rotational velocity in the input to translational velocity in the output).

Combining the drag matrix from (B.16) with the dual adjoint-inverse matrix in (B.18) and the Jacobians from (B.14), then summing over the links, produces a linear map from the system body and shape velocity to the force acting on the body, which we denote ω:

$$F_1^b = \Big(\overbrace{\sum_i Ad_{g_{\lambda_i}}^{-*} \left(-D\right) J_i}^{\omega(\alpha)} \Big) \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}, \tag{B.19}$$

in which the dependence of ω on α is inherited from the J and Ad^{-*} terms.

At low Reynolds numbers, swimmers are at quasistatic equilibrium, with the net external force equal to zero in all directions. This means that the ω term from (B.19) acts as a *Pfaffian constraint* on the feasible velocities for the system: for a set of $\mathring{g}_1$ and $\dot{\alpha}$ velocities to be a solution to the system equations of motion, they must be in the null space of ω:

$$[0] = \omega(\alpha) \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}. \tag{B.20}$$

We can convert this null-space condition on achievable velocities to a mapping from specified $\dot{\alpha}$ shape velocities to their complementary body velocities $\mathring{g}_1$ inside the null space. We first separate ω into two blocks, one of which operates on the body velocity and the second of which operates on the shape velocity,

$$[0] = \begin{bmatrix} \omega_g & \omega_\alpha \end{bmatrix} \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}, \tag{B.21}$$

Taking advantage of the left-hand side of this equation being zero, we then manipulate the blocks of the constraint equation as

$$-\omega_g \mathring{g}_1 = \omega_\alpha \dot{\alpha} \tag{B.22}$$

$$\mathring{g}_1 = -\omega_g^{-1} \omega_\alpha \dot{\alpha} \tag{B.23}$$

$$\mathring{g}_1 = -\mathbf{A}(\alpha)\dot{\alpha}, \tag{B.24}$$

with $\mathbf{A}$ in the final expression serving as a linear map from shape to body velocities of the system, and known in the geometric mechanics literature as the system's *local connection*.

Because $\mathbf{A}$ linearly maps $\dot{\alpha}$ to $\mathring{g}_1$, we can use it to construct a set of Jacobians J_i^α mapping from $\dot{\alpha}$ to the body velocities of the links, with

$$\mathring{g}_i = J_i(\alpha) \begin{bmatrix} -\mathbf{A}(\alpha)\dot{\alpha} \\ \dot{\alpha} \end{bmatrix} = \overbrace{\left(J_i(\alpha) \begin{bmatrix} -\mathbf{A}(\alpha) \\ \mathrm{Id}^{m\times m} \end{bmatrix} \right)}^{J_i^\alpha(\alpha)} \dot{\alpha}, \tag{B.25}$$

where Id is an identity matrix.

In addition to acting as a linear map from body velocity to body force, the drag matrix D from (B.16) also serves as a quadratic map from body velocity of a link to the power required to maintain that velocity in the presence of viscous drag,

$$P_i = F_i^b \cdot \mathring{g}_i = (D\mathring{g}_i)^T \mathring{g}_i = \mathring{g}_i^T D\, \mathring{g}_i. \tag{B.26}$$

Pre- and post-multiplying the drag matrix by the Jacobian in (B.25) pulls it back to an effective drag matrix for that link on the joint angle space,

$$P_i = \dot{\alpha}^T \Big(\overbrace{J_i^{\alpha,T} D\, J_i^\alpha}^{D_i^\alpha} \Big) \dot{\alpha}. \tag{B.27}$$

Summing the D_i^α matrices across the links produces a total drag matrix D^α on the joint angle space,

$$P = \sum_i P_i = \dot{\alpha}^T \Big(\overbrace{\sum_i J_i^{\alpha,T} D\, J_i^\alpha}^{D^\alpha} \Big) \dot{\alpha}. \tag{B.28}$$

As discussed in [35] and the main text, this drag matrix serves as a Riemannian metric on the joint space, such that pathlengths

$$s = \int \sqrt{d\alpha^T D^\alpha \, d\alpha} \tag{B.29}$$

are equal to the time required to traverse paths in the shape space at constant unit power dissipation.

B.5 Subriemannian geometry

From a subriemannian perspective, the local connection and metric tensors from (B.24) and (B.28) can be calculated by first using the link Jacobians J_i to pull the drag matrices back to a metric tensor M on the full configuration space,

$$P = \begin{bmatrix} \overset{\circ}{g}_1^T & \dot{\alpha}^T \end{bmatrix} \Big(\overbrace{\sum_i J_i^T D_i \, J_i}^{M} \Big) \begin{bmatrix} \overset{\circ}{g}_1 \\ \\ \dot{\alpha} \end{bmatrix}. \tag{B.30}$$

We then generate our connection on the configuration space by taking the vertical space of motions as the pure body velocities, and the horizontal space of motions (allowable under the system constraints) as being orthogonal to the vertical velocities with respect to M, i.e.

$$V \subset TQ \ni \begin{bmatrix} \overset{\circ}{g}_1 \\ \\ \mathbf{0} \end{bmatrix} \tag{B.31}$$

$$H \subset TQ \ni v^T M h = 0. \tag{B.32}$$

Taking the metric tensor M as having a block structure

$$M = \begin{bmatrix} M_g & M_{g\alpha} \\ M_{g\alpha}^T & M_\alpha \end{bmatrix} \tag{B.33}$$

(with separations corresponding to the separation between $\mathring{g}_1$ and $\dot{\alpha}$ in the generalized velocity vector), the top section of M must have same nullspace as the ω constraint calculated in (B.19) (and here happens to be equal to ω), because horizontal velocities must produce zero metric product with vertical velocities: the condition

$$\mathbf{0} = \begin{bmatrix} \mathring{g}_1^T & \mathbf{0} \end{bmatrix}_V M \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}_H \tag{B.34}$$

implies that

$$\mathbf{0} = \begin{bmatrix} M_g & M_{g\alpha} \end{bmatrix} \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}_H \tag{B.35}$$

from which we can construct the local connection as $\mathbf{A} = M_g^{-1} M_{g\alpha}$.

Once we have constructed this local connection, we can pull back the metric M from the full configuration space to the base space as

$$P = \dot{\alpha}^T \Big(\overbrace{\begin{bmatrix} -\mathbf{A}^T & \mathrm{Id}^{m\times m} \end{bmatrix} M \begin{bmatrix} -\mathbf{A} \\ \mathrm{Id}^{m\times m} \end{bmatrix}}^{M_B} \Big) \dot{\alpha}, \tag{B.36}$$

with the result that M_B is equal to the D^α drag matrix that we calculated in (B.28).

B.6 Changing frames

For visualizing system motion, it is often helpful to use a base frame that is not the first link. For example, many system symmetries are more apparent if we use the middle link of a chain, or a generalized center-of-mass frame at an averaged position and orientation of the links. If we designate the new base frame as link 0, we can take the position of the new body frame relative to the original link as $g_{\frac{0}{\lambda_1}}$, and the positions of the other links relative to the new body frame are

$$g_{\frac{\lambda_i}{0}} = (g^{-1}_{\frac{0}{1}})(g_{\lambda_i}). \tag{B.37}$$

Similarly, we can use the position of the new body frame relative to the original base link to transform the link Jacobians so that they take the body velocity of the new frame and the shape velocities as inputs,

$$\overset{\circ}{g}_i = J_{\frac{i}{0}} \begin{bmatrix} \overset{\circ}{g}_0 \\ \dot{\alpha} \end{bmatrix}. \tag{B.38}$$

When calculating $J_{\frac{i}{0}}$, we need to account for the fact that changing the joint angles moves the new body frame relative to the original base link. To make the calculation, we first start by finding the Jacobian from the joint angular velocities to the new frame's body velocity, with the original body frame held fixed. This is

the derivative of the g_0 frame's position with respect to the joint angles, rotated into the g_0 frame as

$$J_{\frac{0}{\alpha}} = \begin{bmatrix} \cos\theta_0 & \sin\theta_0 & 0 \\ -\sin\theta_0 & \cos\theta_0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \frac{\partial g_0(\alpha)}{\partial \alpha}. \tag{B.39}$$

Once we have found this Jacobian, we can separate the original link Jacobians into the blocks that interact with base-link and joint motion,

$$\mathring{g}_i = \left[J_{i,g} \;\middle|\; J_{i,\alpha} \right] \begin{bmatrix} \mathring{g}_1 \\ \dot{\alpha} \end{bmatrix}, \tag{B.40}$$

and use these blocks to calculate the new link velocities as

$$\mathring{g}_i = \overbrace{\left[(J_{i,g}\ Ad^{-1}_{g_{\frac{\lambda_1}{0}}}) \;\middle|\; (J_{i,\alpha} - Ad^{-1}_{g_{\frac{\lambda_i}{0}}} J_0) \right]}^{J_{\frac{i}{0}}} \begin{bmatrix} \mathring{g}_0 \\ \dot{\alpha} \end{bmatrix}, \tag{B.41}$$

where the $Ad^{-1}_{g_{\frac{\lambda_1}{0}}}$ term in the first block maps $\mathring{g}_0$ to $\mathring{g}_1$, on which $J_{i,g}$ then acts to produce its contribution to $\mathring{g}_i$. The $-Ad^{-1}_{g_{\frac{\lambda_i}{0}}} J_0$ transformation on the second block makes its contribution to the link body velocity relative to the new body frame's motion, rather than relative to the original base link's motion. Alternatively, the

new link Jacobians can also be written directly as

$$\mathring{g}_i = \overbrace{\left[Ad^{-1}_{g_{\frac{\lambda_i}{0}}} \;\middle|\; J_{i,\alpha} - Ad^{-1}_{g_{\frac{\lambda_i}{0}}} J_0 \right]}^{J_{\frac{i}{0}}} \begin{bmatrix} \mathring{g}_0 \\ \dot{\alpha} \end{bmatrix}, \tag{B.42}$$

in which the first block (corresponding to rigid motion of the system) is recognized as being simply the adjoint mapping associated with the position of the link relative to the new body frame.

Once the $g_{\frac{\lambda_i}{0}}$ and $J_{\frac{i}{0}}$ terms have been calculated, they may used wherever the g_{λ_i} and J_i terms were previously used, such as in the calculations of low Reynolds numbers described above.

B.7 Continuous backbones

To calculate the kinematics of the continuous-backbone systems, we use an essentially-similar procedure to that for the discrete-link system, except that the product and summing operations are replaced by integrals. In place of discrete joint angles, the shape of a continuous system is given by its curvature κ, which is the rate at which the tangent direction of the body changes along the length of the body. Taking κ as a function of the shape variables α and the position s along the body, the displacement of a frame tangent to the body at point s relative to the tangent frame at $s = 0$ is the integral along the body of a vector that flows along the body

at unit speed while rotating at rate κ,

$$h(\alpha, s) = \int_0^s \begin{bmatrix} \cos\theta(\alpha, S) & \sin\theta(\alpha, S) & 0 \\ -\sin\theta(\alpha, S) & \cos\theta(\alpha, S) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \kappa(\alpha, S) \end{bmatrix} dS. \quad \text{(B.43)}$$

The Jacobian from shape velocity to body velocity of a tangent-frame on the system is the gradient of $h(\alpha, s)$ with respect to α, rotated into that tangent frame,

$$J_\alpha(\alpha, s) = \begin{bmatrix} \cos\theta(s) & \sin\theta(s) & 0 \\ -\sin\theta(s) & \cos\theta(s) & 0 \\ 0 & 0 & 1 \end{bmatrix} \nabla_\alpha h(\alpha, s), \quad \text{(B.44)}$$

where by taking the gradient of (B.43), making an exchange in the order of integration and derivation, and applying the product rule, the gradient of h is

$$\nabla_\alpha h(\alpha, s) = \int_0^s \left(\nabla_\alpha \theta \begin{bmatrix} -\sin\theta & -\cos\theta & 0 \\ \cos\theta & -\sin\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \kappa \end{bmatrix} + \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \nabla_\alpha \kappa \end{bmatrix} \right) dS, \quad \text{(B.45)}$$

in which (by a second exchange of integration and derivation, this time applied to the θ component of (B.43))

$$\nabla_\alpha \theta(\alpha, s) = \int_0^s \nabla_\alpha \kappa(\alpha, S)\, dS. \tag{B.46}$$

Once $h(\alpha, s)$ and $J_\alpha(\alpha, s)$ have been found, $J(\alpha, s)$ can be constructed by concatenating it with the adjoint inverse from the base frame to the tangent frame at s,

$$J(\alpha, s) = \left[Ad^{-1}_{h(\alpha,s)} \;\middle|\; J_\alpha(\alpha, s) \right], \tag{B.47}$$

enabling computation of the Pfaffian, shape-space drag metric, and full-configuration drag metric by substituting $\int_0^s$ for $\sum_i$ into their respective equations.

Appendix C: Comparison with Lighthill Efficiency

Lighthill's efficiency is defined as the mechanical power required to drag the swimmer at a given velocity, v_{req}, at a fixed configuration, by an externally applied force, divided by the actual power required for swimming at the same velocity,

$$\eta_{\text{LH}} = \frac{P_{ref}}{P_{avg}}, \tag{C.1}$$

where P_{ref} is the mechanical power required to drag the swimmer at a given velocity v_{req} at a fixed configuration, and P_{avg} is the average power required for the swimmer swimming at the same velocity. Note that, P_{avg} is a function of the gait the swimmer is executing, ϕ, and P_{ref} is constant across all gaits.

It is readily shown that for a given swimmer, η_{LH} is proportional to $\frac{v_{avg}^2}{P_{avg}}$ [56]. Ignoring the proportionality factor we can write,

$$\eta_{\text{LH}} = \frac{v_{avg}^2}{P_{avg}} \tag{C.2}$$

where v_{avg} is the average velocity of the swimmer.

In this paper we use a geometric measure of efficiency, η_2, defined in (4.18) as:

$$\eta_2 = \frac{g_\phi}{s}. \tag{C.3}$$

In this appendix, we show that for fully actuated swimmers this definition is equivalent to Lighthill's efficiency, i.e. we will show that the gait that maximizes η_2, maximizes Lighthill's efficiency and vice versa.

Since s is the pathlength of the gait in the shape space as measured under the power metric, we can rewrite η_2 as:

$$\eta_2 = \frac{g_\phi}{\int_0^T \sqrt{P(t)}dt} \tag{C.4}$$

From [7], we know that the optimal pacing for any gait utilizes a constant power pacing. Hence for all t,

$$P(t) = P_{avg} \tag{C.5}$$

Substituting (C.5) into the expression for η_2, we get,

$$\begin{aligned} \eta_2 &= \frac{g_\phi}{\sqrt{P_{avg}}T} && \text{(C.6)} \\ &= \frac{v_{avg}}{\sqrt{P_{avg}}} && \text{(C.7)} \\ &= \sqrt{\eta_{\text{LH}}} && \text{(C.8)} \end{aligned}$$

Since square root is a monotonic function, a gait that maximizes our definition of efficiency also maximizes the Lighthill efficiency and vice versa.

Appendix D: Local Commutativity Assumption

In this subsecton, we quantify the errors introduced by the local commutativity assumption in obtaining the Lie bracket formula (5.12). In this analysis, we will restrict ourselves to square gaits of side length $2L$ and we start the gait from the bottom left corner. We also assume a linearly varying vector field. Note that, while we restrict our analysis to square shaped gaits, we can always remap any closed loop in the shape space onto a square through a change of coordinates for the shape space. From the local commutativity assumption, for the first quarter of the gait we get

$$g^{-1}(0)g(\tfrac{T}{4}) = \prod_{t=0}^{T/4} \exp(\mathbf{A}(t)dt) \approx \exp\left(\int_0^{T/4} dt_1 \mathbf{A}(t_1)\right) \tag{D.1}$$

Since we have assumed a linearly varying vector field, we can use Taylor series expansion to obtain the value of the connection vector field over the shape space,

$$\begin{aligned} g^{-1}(0)g(\tfrac{T}{4}) &\approx \exp\Big(\int_0^{\frac{T}{4}} dt_1 \mathbf{A}(t_1)\Big) && \text{(D.2)} \\ &= \exp\Big(\int_{(-L,-L)}^{(L,-L)} \mathbf{A}(r_1, r_2)dr\Big) && \text{(D.3)} \\ &= \exp\Big(\int_{(-L,-L)}^{(L,-L)} \Big(\mathbf{A}|_{(0,0)} + r_1\frac{\partial \mathbf{A}}{\partial \alpha_1} + r_2\frac{\partial \mathbf{A}}{\partial \alpha_2}\Big)dr\Big) && \text{(D.4)} \\ &= \exp\Big(2L\Big(\mathbf{A}_1|_{(0,0)} - L\frac{\partial \mathbf{A}_1}{\partial \alpha_2}\Big)\Big) && \text{(D.5)} \end{aligned}$$

We can obtain similar expressions for displacements resulting from other quarters of the gait:

$$g^{-1}(\tfrac{T}{4})g(\tfrac{T}{2}) \approx \exp\Big(2L\Big(\mathbf{A}_2|_{(0,0)} + L\frac{\partial \mathbf{A}_2}{\partial \alpha_1}\Big)\Big) \tag{D.6}$$

$$g^{-1}(\tfrac{T}{2})g(\tfrac{3T}{4}) \approx \exp\Big(-2L\Big(\mathbf{A}_1|_{(0,0)} - L\frac{\partial \mathbf{A}_1}{\partial \alpha_2}\Big)\Big) \tag{D.7}$$

$$g^{-1}(\tfrac{3T}{4})g(T) \approx \exp\Big(-2L\Big(\mathbf{A}_2|_{(0,0)} + L\frac{\partial \mathbf{A}_2}{\partial \alpha_1}\Big)\Big) \tag{D.8}$$

To make the equations describing displacement resulting from each quarter of the gait to the magnitude of the connection vector fields and size of the gait concise, we define new variables:

$$\alpha = 2L\mathbf{A}_1|_{(0,0)} \tag{D.9}$$

$$\beta = 2L\mathbf{A}_2|_{(0,0)} \tag{D.10}$$

$$\delta = 4L^2\frac{\partial \mathbf{A}_1}{\partial \alpha_2} \tag{D.11}$$

$$\Delta = 4L^2\frac{\partial \mathbf{A}_2}{\partial \alpha_1} \tag{D.12}$$

Substituting these variables into (D.5)-(D.8), for the first quarter of the gait, we get

$$g^{-1}\Big(0\Big)g\Big(\frac{T}{4}\Big) \approx \exp\Big(\int_0^{\frac{T}{4}} dt_1\mathbf{A}(t_1)\Big) = \exp(\alpha - \delta/2) \tag{D.13}$$

Similarly when we divide the entire gait into 8 equal parts, we can show that

upto $O(L^3)$

$$\begin{aligned} g^{-1}\Big(0\Big)g\Big(\frac{T}{4}\Big) &= g^{-1}\Big(0\Big)g\Big(\frac{T}{8}\Big)g^{-1}\Big(\frac{T}{8}\Big)g\Big(\frac{T}{4}\Big) && \text{(D.14)} \\ &= \exp\Big(\frac{\alpha}{2}-\frac{\delta}{4}+\frac{\delta_1}{2}\Big)\exp\Big(\frac{\alpha}{2}-\frac{\delta}{4}-\frac{\delta_1}{2}\Big) && \text{(D.15)} \\ &\approx \exp\Big(\alpha-\frac{\delta}{2}+\frac{1}{4}[\delta_1,\alpha]\Big) && \text{(D.16)} \end{aligned}$$

where,

$$\delta_1 = -L^2\frac{\partial A_1}{\partial \alpha_1} \qquad \text{(D.17)}$$

Dividing the gait into 2^{k+2} parts yields,

$$\begin{aligned} g^{-1}\Big(0\Big)g\Big(\frac{T}{4}\Big) &= g^{-1}\Big(0\Big)g\Big(\frac{T}{2^{k+2}}\Big)\cdots g^{-1}\Big(\frac{(2^k-1)T}{2^{k+2}}\Big)g\Big(\frac{T}{4}\Big) && \text{(D.18)} \\ &= \exp\Big(\alpha-\frac{\delta}{2}+(\frac{1}{4}+\frac{1}{16}+\cdots+\frac{1}{4^k})[\delta_1,\alpha]\Big)\Big) && \text{(D.19)} \end{aligned}$$

Thus, in the limit of k approaching ∞, we get,

$$g^{-1}\Big(0\Big)g\Big(\frac{T}{4}\Big) = \exp\Big(\alpha-\frac{\delta}{2}+\frac{1}{3}[\delta_1,\alpha]\Big) \qquad \text{(D.20)}$$

Appendix E: Starting Point Selection Heuristic

In this appendix, we show that a good heuristic for selecting starting points that reduce leading order error is to select a point such that the tangent to the gait at that point coincides with the general direction of rotational vector field flow in the region enclosed by the gait. We assume the system has a constant rotational vector field to obtain this heuristic.

Since the rotational connection vector field is constant, we know that there is no net rotation produced by any gait, i.e. for any gait ϕ, the net rotation produced on executing the gait, g_ϕ^θ is

$$g_\phi^\theta = \iint D\mathbf{A}^\theta = 0. \tag{E.1}$$

Recall that the leading order error term associated with a gait is given by (5.53) as:

$$E = \frac{1}{2}\left[(\alpha+\beta), \iint D\mathbf{A}\right]. \tag{E.2}$$

where

$$\alpha+\beta = \frac{(a+b)-(c+d)}{2} = \frac{1}{2}\left(\int_0^{\frac{T}{2}} dt_1 \mathbf{A}(t_1) - \int_{\frac{T}{2}}^{T} dt_1 \mathbf{A}(t_1)\right). \tag{E.3}$$

So the leading order error is equal to the Lie bracket of difference in the line integral

of the connection vector field over the first and second half of the gait with the displacement predicted by the total Lie bracket formula. The Lie bracket of two elements $a, b \in se(2)$ is given by

$$\left[\begin{pmatrix} a^x \\ a^y \\ a^\theta \end{pmatrix}, \begin{pmatrix} b^x \\ b^y \\ b^\theta \end{pmatrix}\right] = \begin{pmatrix} b^\theta a^y - a^\theta b^y \\ a^\theta b^x - b^\theta a^x \\ 0 \end{pmatrix} \tag{E.4}$$

Because of this structure, we can use the fact that no net rotation is produced on executing a gait to simplify the leading order error expression to

$$E = \frac{1}{2}\left[(\alpha+\beta), \iint D\mathbf{A}\right] = \frac{1}{2}(\alpha+\beta)^\theta \begin{pmatrix} -\iint D\mathbf{A}^y \\ \iint D\mathbf{A}^x \\ 0 \end{pmatrix}. \tag{E.5}$$

Thus for any image family of gaits, an optimal starting point that minimizes the leading order error, E, would be one where the θ component of $\alpha+\beta$ is zero, i.e.

$$(\alpha+\beta)^\theta = \frac{1}{2}\left(\int_0^{\frac{T}{2}} dt_1 \mathbf{A}^\theta(t_1) - \int_{\frac{T}{2}}^{T} dt_1 \mathbf{A}^\theta(t_1)\right) = 0. \tag{E.6}$$

Since the vector fields are constant, we also know that no net rotation is produced when executing a gait, i.e.,

$$g_\phi^\theta = \left(\int_0^{\frac{T}{2}} dt_1 \mathbf{A}^\theta(t_1) + \int_{\frac{T}{2}}^{T} dt_1 \mathbf{A}^\theta(t_1)\right) = 0 \tag{E.7}$$

Solving the system of equations described by (E.6) and (E.7), we get that the leading order error E is minimized when

$$\left(\int_0^{\frac{T}{2}} dt_1 \mathbf{A}^{\theta}(t_1) = \int_{\frac{T}{2}}^{T} dt_1 \mathbf{A}^{\theta}(t_1) \right) = 0. \tag{E.8}$$

In other words, for a system with constant rotational vector field such as the differential drive car, the leading order error associated with any gait is minimized when each half of the gait produces no rotation. Similarly, it is also possible to show that starting points which maximize the leading order error E are ones which maximize the magnitude of rotation produced by each half of the gait. For a system with a constant rotational field, this is equivalent to selecting a point such that the tangent to the gait at the point coincides with the direction of the rotational vector field.