AN ABSTRACT OF THE THESIS OF

Capprin James Bass for the degree of Master of Science in Robotics presented on June 8, 2022.

Title: Geometric Optimization Methods for Mobile Systems

Abstract approved: _____

Ross L. Hatton

When studying robot systems, it is common to ask about optimal approaches to accomplish a given task. In the context of mobile systems, particularly biomimetic systems, optimization tasks are closely related to the relevant dynamics of locomotion. In this thesis, building on prior work from the geometric mechanics community, we introduce several methods for optimization in the context of mobile systems. Each topic is covered in a separate manuscript: in the first manuscript, we investigate how error accrues when using second order displacement approximations for planar mobile systems, and how coordinate optimization reduces error by accounting for higher order error effects. In the second manuscript, we extend planar coordinate optimization to the group of 3D rotations, clarifying methods presented in prior work. Contributions introduced here further reduce error inherent to displacement approximations. In the third and final manuscript, we use recent robot gait optimization methods to perform dimensionality reduction on robot shape spaces, respecting useful dynamics. The resulting 2D shape spaces have useful topological and geometric properties, and could act as a building block for more general robot dimensionality reduction. Finally, the conclusion summarizes the shared themes of the included papers, and discusses in general the potential of geometric mechanics in future robotics work.

Geometric Optimization Methods for Mobile Systems

by

Capprin James Bass

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Master of Science

Presented June 8, 2022
Commencement June 2022

Master of Science thesis of Capprin James Bass presented on June 8, 2022

APPROVED:

Major Professor, representing Robotics

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.

Capprin James Bass, Author

ACKNOWLEDGEMENTS

Academic

I am indebted to my advisor, Dr. Ross L. Hatton. Your counsel has elevated my understanding of mathematics and altered my career trajectory. It has been a privilege to work with you.

Personal

I wish to thank all the members of my family. This body of work is a consequence of your care and attention throughout my life, and my accomplishments are as much yours as they are my own.

CONRIBUTION OF AUTHORS

Suresh Ramasamy initiated research into higher order terms of the Baker-Campbell-Haudorff series; in particular, gait amplitude and phase relationships were first discussed in his own dissertation. His analysis formed the foundational ideas of the manuscript included here.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

TABLE OF CONTENTS (Continued)

LIST OF FIGURES

GEOMETRIC OPTIMIZATION METHODS FOR MOBILE SYSTEMS

1. GENERAL INTRODUCTION TO GEOMETRIC MECHANICS

This thesis is a study of mobile robot systems, particularly those behaving similarly to biological organisms. Biomimetic robots depend heavily on specific dynamic interactions with their environment; control policies must balance a broad set of physics in order to produce effective locomotion. With this in mind, a key area of interest is *optimal* design and control of mobile systems. Our objective is to understand the best way to balance the given dynamics of a system with specific locomotion objectives. Often, we also want to build intuition about how a system behaves so that design decisions in the future are as natural as possible. In general, our approach to these problems relies on the relationship between a system's *geometry* and its underlying *mechanics*.

The set of methods in this thesis builds on a broader set of approaches by the geometric mechanics community. The term "geometric mechanics" refers to an application of differential geometry to mechanics. The objects studied in differential geometry often have a natural correspondence with the sorts of structures we see in mechanics. Notions of *form* (length, area, volume...) and *change* (in the form of derivatives) are baked into differential geometry. These same concepts are relevant to mechanics. For example, kinematic quantities (position and its derivatives) occupy manifolds with curvature, tangent spaces, and tangent-tangent spaces; dynamic quantities (force, torque) are covectors over tangent spaces. By understanding the kinds of mathematical objects we are working with, certain operations become much more natural than they are when expressed in other ways.

This thesis introduces a number of these geometric approaches to mechanics, and explains further why these tools are useful.

The body of this thesis is organized into three separate manuscripts, all sharing a common theme. In the first manuscript, *Characterizing Error in Noncommutative Geometric Gait Analysis*, we investigate the extent that coordinate optimization reduces error when estimating displacement for mobile systems. We construct a bound on third order displacement errors, and demonstrate that optimized minimum perturbation coordinates significantly reduce error. In the second manuscript, *Minimum Perturbation Coordinates on the 3D Rotation Group*, we extend our own prior work and compute minimum perturbation coordinates on $SO(3)$, the group of 3D rotations. We also introduce additional mathematics that improve coordinate optimization results. In the third and final manuscript, *Geometric Nonlinear Dimensionality Reduction for Robot Shape Spaces*, we identify a "most useful" 2D shape space for systems with many degrees of freedom, using prior tools from the geometric mechanics community. This work is a robust proof-of-concept, and may be extended as a general dimensionality reduction for mobile systems.

The shared theme of this work is an investigation of optimization methods for mobile systems: optimizing coordinate choice in the first two manuscripts, and optimizing control domain in the final manuscript. However, across this thesis, we explore the tools required to use geometric mechanics on more real-world systems, including those that move spatially (in 3D space) or have many degrees of freedom. With more intuitive understanding of how these systems behave mechanically, we can construct and control agents that are prepared to interact with and inhabit the natural world.

Characterizing Error in Noncommutative Geometric Gait Analysis

Capprin Bass, Suresh Ramasamy, and Ross L. Hatton

Abstract

A key problem in robotic locomotion is in finding optimal shape changes to effectively displace systems through the world. Variational techniques for gait optimization require estimates of body displacement per gait cycle; however, these estimates introduce error due to unincluded high order terms. In this paper, we formulate existing estimates for displacement, and describe the contribution of low order terms to these estimates. We additionally describe the magnitude of higher (third) order effects, and identify that choice of body coordinate, gait diameter, and starting phase influence these effects. We demonstrate that variation of such parameters on two example systems (the differential drive car and Purcell swimmer) effectively manages third order contributions.

Proceedings of the 2022 IEEE International Conference for Robotics and Automation

2.1. Introduction

In nature, creatures of all kinds move with gaits. Bodies interact with their environment through changes in shape, which displace the body through the world. These gaits are, by nature, cyclic: running, swimming, and flying all involve repeated action, and result in body displacement. It is useful to describe the locomotion of robots in the same way.

When controlling locomoting robots, it is useful to understand which gait cycles result in "good" displacements, based on desired gait properties such as displacement per unit time or unit energy [1,2]. As one approach to this problem, the geometric mechanics community has described a framework for relating system dynamics, configuration, and gait geometry that provides insight into the displacements resulting from particular gaits [3–11].

Because gaits are cycles in system shape, they form closed loops in the shape space of the system. The *motility map* $\mathbf{A}$, defined over the shape space of a system, can be used to map shape velocity to body velocity [8]. Using the corrected body velocity integral (cBVI), a surface integral of the *total Lie bracket* over the region enclosed by a gait, we construct an estimate for displacement, g_ϕ [1,12],

$$g_\phi \approx \exp\left(\underbrace{\iint_\phi \overbrace{d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]}^{\text{total Lie bracket}}}_{\text{cBVI}}\right), \tag{2.1}$$

in which the first term ($d\mathbf{A}$) captures the nonconservativity of locomotion; this is the "forwards minus backwards" displacement due to the gait. The second term ($[\mathbf{A}_1, \mathbf{A}_2]$) is the local Lie bracket of the matrix columns of the motility map,[1] and captures the effects of noncommutativity of the position space; this is the sideways "parallel parking"

[1]The local Lie bracket may be extended to greater than two dimensions by taking Lie brackets of each matrix column: $\sum_{j>i}[\mathbf{A}_i, \mathbf{A}_j]$ [1].

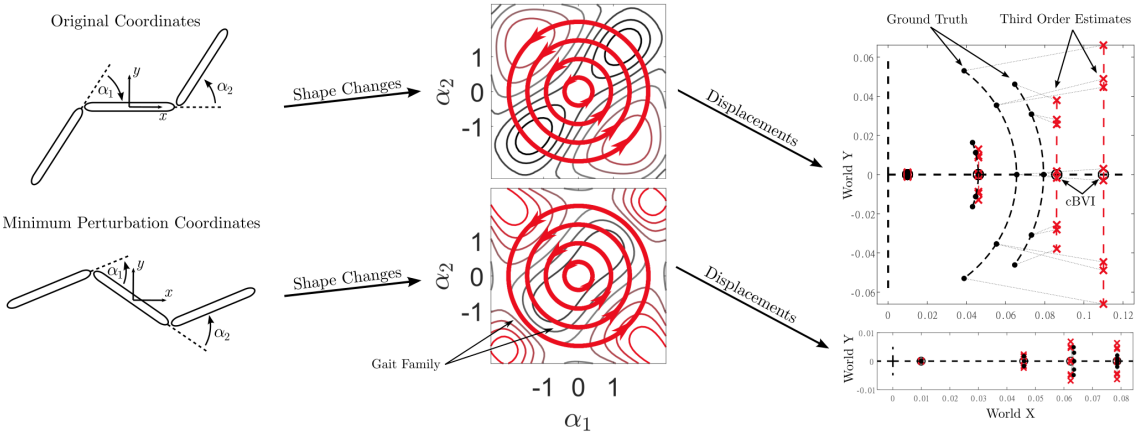

FIGURE 2.1: Relationship between systems and displacements. The configuration of a system is represented by shape variables; in this case, the Purcell swimmer is represented by two relative orientations α_1 and α_2. A *gait family* captures shape changes of a certain kind. Integration of the BCH series approximation inside the gait provides estimates for displacement. As gait diameter increases, so does displacement per cycle; starting phase impacts gait error angle.

effect from "move forward and turn" actions. Fig. 2.1 captures this relationship between system, shape changes, and estimated displacements for the Purcell swimmer [13].

Our previous work has shown that the cBVI has associated error; that is, the displacement predicted by the cBVI is not exactly ground truth. This error comes from unaccounted-for higher order displacement effects present in many systems and gaits. We have argued that particular choices of body coordinates (in particular, the use of minimum perturbation coordinates [14]) reduce the contribution of higher order terms, instead capturing their effect with the total Lie bracket [12]. However, we have not previously quantified the error introduced by these higher order terms.

In this paper, we address the specific gap of understanding in the magnitude and direction of higher order terms of the total Lie bracket. We do so by using the Baker-Campbell-Hausdorff series to construct an expression for the cBVI that includes higher

order terms:

$$g_\phi = \exp\left(\iint_\phi \overbrace{\left(d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]\right)}^{\text{total Lie bracket}} + \underbrace{\frac{\pi\ell}{8}\left[\bar{\mathbf{A}}, \iint_\phi \left(d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]\right)\right]}_{\text{third order effects}} + \cdots \right), \tag{2.2}$$

in which $\bar{\mathbf{A}}$ is an estimate for the average of the motility map in the region of the gait and ℓ is the characteristic diameter of the gait in the shape space.

We comment further on the factors contributing to leading order error. Third order effects (in the plane) are bounded by

$$||\mathbf{A}|| \cdot ||D\mathbf{A}||\ell^3, \tag{2.3}$$

where $D\mathbf{A} = d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]$, referring to to the total Lie bracket. Because the cBVI is an area integral of $D\mathbf{A}$, third order effects may be expressed *relatively* to the cBVI as being proportional to

$$||\mathbf{A}||\ell, \tag{2.4}$$

the magnitude of the motility map and the size of the gait. Given $||\mathbf{A}||$ for a system at a given point, (2.4) communicates the maximum size gait possible before error becomes too large. Coordinate choice also affects $||\mathbf{A}||$; in fact, our own minimum perturbation coordinates maximize ℓ for a given level of acceptable error. Fig. 2.1 demonstrates this effect for two choices of coordinate and several gait amplitudes.

In addition, for gaits without net body rotation, third order effects are directed orthogonally to the displacement predicted by the cBVI. As a result, their respective contributions may either be compared in terms of absolute magnitude or in terms of the "error angle" resulting from third order effects.

The rest of the paper is organized as follows: In §2.2., we describe the model, as well as the supporting mathematics leading up to the total Lie bracket. In §2.3., we construct

the total Lie bracket. In §2.4., we approximate third order effects, and derive expressions for third order contributions and a heuristic on characteristic gait diameter. In §2.5., we apply these methods to two locomoting systems. In §2.6., we make concluding remarks and comment on future work.

2.2. Model Background

2.2.1 Model

We model our systems as having a configuration space Q, partitioned into a position space G and a shape space R, as in [2]. Elements $g \in G$ describe positions of the system in space, and $r \in R$ describe the shape of the system itself. Fig. 2.1 illustrates the difference between position and shape. As in [2,8,12,14], the *local connection* (or motility map) may be used to map infinitesimal shape changes to infinitesimal position changes,[2]

$$\mathring{g} = \mathbf{A}(r)\dot{r}, \tag{2.5}$$

in which $\mathbf{A}$ refers to the local connection,[3] and $\mathring{g}$ is a body velocity. Body velocities are elements of the Lie algebra of the position space; they represent velocity in the local frame. As such, Lie algebra elements may be represented by either a column vector (with the body frame acting as bases) or in a corresponding matrix form. For the remainder of this paper, the position space is the special Euclidean group ($g \in SE(2)$); we notate body velocities with $\mathring{g} \in \mathfrak{se}(2)$.

2.2.2 Gaits

Certain changes in system shape result in a displacement through the position space. In the context of locomotion, it is useful to describe shape changes in terms of cyclic *gaits*,

[2]This expression makes the assumption that systems behave *kinematically*. Previous work [15] extends this domain to apply to many systems.

[3]In previous work, the local connection, by convention, encodes negative body motion; we have dropped this convention for this paper.

where a mapping $\phi : [0, T) \to R$ describes the shape r at time $t \in [0, T)$; T is the period of the gait. This structure allows us to express displacement from the identity induced by a gait:

$$g_\phi = \int_0^T g(t)\mathbf{A}(r(t))\dot{r}(t)dt = \oint_\phi g\mathbf{A}(r)dr, \tag{2.6}$$

where the rightmost integral described in (2.6) is a path integral along a closed loop drawn in the shape space by the gait ϕ.

This integral is invariant to time parameterizations, but does depend on the ordering of actions along the path. Both versions of the integral contain system configuration, which a traditional Riemann integral does not adequately describe. In contrast, the product-integral[4] accounts for the ordering of actions along the path, respecting configuration:

$$g_\phi = \prod_0^T \Big(\exp\big(\mathbf{A}(t)\dot{r}(t)dt\big)\Big), \tag{2.7}$$

in which exponentiating $\mathbf{A}(t)\dot{r}(t)$ over infinitesimal time produces the corresponding body frame transformation, and taking the product of all these infinitesimal transformations produces the total displacement over the gait.

Because multiplication of translation/rotation elements does not commute, we still cannot compute a closed form expression for this integral. By employing the Baker-Campbell-Hausdorff series described in the next section, however, we can construct an approximate closed-form solution that provides geometric insight into the system motion.

2.2.3 The Baker-Campbell-Hausdorff Series

The Baker-Campbell-Hausdorff (BCH) series expresses the result of executing serial group actions as a single equivalent operation. It is related to the exponential map, which implies a correspondence between a groupwise velocity and a group action. For example,

[4]The product-integral is a *multiplicative* version of the additive Riemann integral. Product integration preserves the effect of the group operation, rather than integrating components independently. In effect, the product integral preserves the order that events occur, as in (2.6).

take a groupwise velocity $\mathring{g} \in \mathfrak{se}(2)$; it is mapped to a group element $g \in SE(2)$ with:

$$g = \exp\left(\mathring{g}\right), \tag{2.8}$$

where the exponential map is equivalent to *integration* of the groupwise velocity over unit time.

Now, take two groupwise velocities $X, Y \in \mathfrak{se}(2)$. Applying their corresponding group actions in series has the form

$$g = \exp\left(X\right)\exp\left(Y\right). \tag{2.9}$$

The BCH series can be used to replace the right hand side of (2.9) with the exponential of a single groupwise velocity $Z \in \mathfrak{se}(2)$ defined such that

$$\exp(Z) = \exp(X)\exp(Y). \tag{2.10}$$

The BCH series is infinite, and its lowest order terms are:

$$Z = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X - Y, [X, Y]] + \cdots. \tag{2.11}$$

Note that the BCH series contains the nominal $X + Y$ as expected from commutative algebra; however, it also contains additional, corrective terms. The following example builds intuition for these terms; refer to Fig. 2.2 for its visualization.

Consider the example of a diffdrive car, which can *drive forward* and *turn*. We assign X as driving forward, and Y as turning; because $X, Y \in \mathfrak{se}(2)$, we write each action as

$$X = [\mathring{x} \quad 0 \quad 0]^T, \quad Y = [0 \quad 0 \quad \mathring{\theta}]^T, \tag{2.12}$$

for some $\mathring{x}, \mathring{\theta}$. The composite motion $\exp(X)\exp(Y)$ encodes displacement after driving forward for some time, and then turning. This results in a $(x, 0, \theta)$ position. In contrast, $\exp(X + Y)$ encodes displacement after simultaneously driving forward *and* turning, resulting in a (x, y, θ) position.

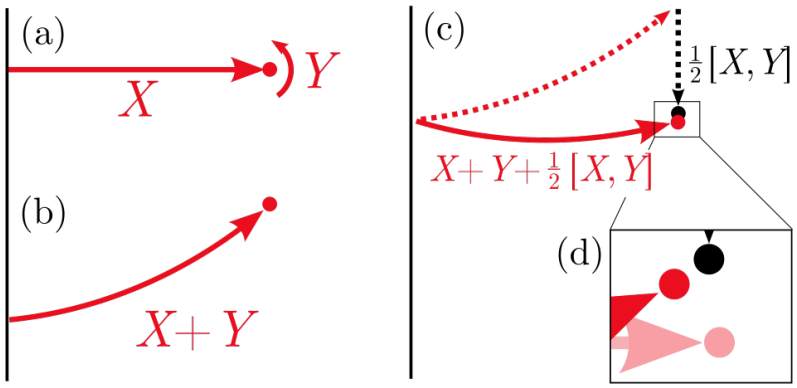

FIGURE 2.2: Action of the BCH series. (a) For the serial actions X, Y, the ground truth is moving forward, followed by a turn. (b) The action $X + Y$ results in movement along an arc. (c) The inclusion of a local Lie bracket term, $\frac{1}{2}[X, Y]$, corrects some error. (d) The endpoints for ground truth (pink), Lie bracket correction (black), and truncated BCH series (red) are distinct.

The two operations result in different predicted displacements of the car in space; $\exp(X)\exp(Y)$ is the ground truth, and $\exp(X+Y)$ is an approximation of ground truth, discarding information about the order in which events occur. To improve the approximation, we can introduce additional terms from the BCH series. The second term expands the approximation to $\exp(X + Y + \frac{1}{2}[X, Y])$. The local Lie bracket captures the fact that X occurred *before* Y, and introduces a lateral velocity to the car that corrects most of the y error in the approximation.

2.3. Baker-Campbell-Hausdorff for Gaits

Rather than integrating (2.6) or (2.7) directly, we construct an integral estimate that captures the relationship between system properties and displacement. This estimate leverages the BCH series to describe leading-order displacement effects from gaits, while simplifying the integral expression such that it may be solved numerically. The total Lie bracket is a truncation of this BCH series expression, as we will show.

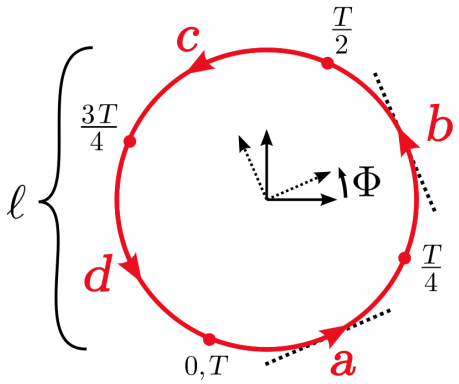

FIGURE 2.3: To construct a displacement integral, a circular gait of diameter ℓ and starting phase Φ is discretized into four segments a, b, c, d.

We first split the gait into four sections a–d, distributed evenly around the gait such that the mean tangent vectors in a and c are antiparallel, as are the mean tangent vectors in b and d, as illustrated in Fig. 2.3. This split discretizes the product-integral from (2.7) into the product of four product-integrals over smaller intervals,

$$g_\phi \approx \prod_{i=1}^{4} \left(\prod_{(i-1)T/4}^{(i)T/4} \Big(\exp\big(\mathbf{A}(t)\dot{r}(t)dt\big) \Big) \right). \tag{2.13}$$

If we apply the BCH series recursively to the (infinite number of infinitesimal) elements in each of the four product integrals and assume commutativity within each gait segment such that all the Lie bracket terms go to zero, the segment integrals may be written as

$$\prod_{\tau} \Big(\exp\big(\mathbf{A}(t)\dot{r}(t)dt\big) \Big) = \exp\left(\int_{\tau} \mathbf{A} d\tau \right), \tag{2.14}$$

where τ is an arbitrary gait segment.[5]

Assigning each of the $\int \mathbf{A}$ integrals to their corresponding segment names, we can rewrite (2.13) as the product of four exponential terms,

$$g_\phi \approx e^a e^b e^c e^d. \tag{2.15}$$

[5]As per [16], the assumption of local commutativity introduces fourth-order errors; we constrain our focus in this paper to third-order errors.

Applying the BCH formula to each term in this product produces a series expression for g_ϕ in terms of the $\int \mathbf{A}$ integrals,

$$\approx \exp(a+b+c+d+\frac{1}{2}([a,b]+[b,c]+\cdots, \tag{2.16}$$

in which the first terms are the "simple integral" of the body frame motions the system makes, and the Lie bracket terms are corrections to the global motion based on the order in which the segments appear in the gait.

With the series expression from (2.16) in hand, we can now use the geometric arrangement of the segments to gain further insight about the gait displacement integral: Because the elements of the gait pairs $\{a,c\}$ and $\{b,d\}$ are antiparallel, we can approximate them in terms of the mean value of the local connection in the region of the gait, its first derivative over the shape space, and the characteristic diameter ℓ of the gait in the shape space. This approximation takes the form

$$\begin{aligned} a &\approx \alpha - \delta/2, & b &\approx \beta + \Delta/2, \\ c &\approx -(\alpha+\delta/2), & d &\approx -(\beta-\Delta/2), \end{aligned} \tag{2.17}$$

in which α and β are the mean values of the columns of the local connection in directions aligned with the a and b sections of the gait, scaled by $\ell\pi/4$, the length of a quarter-circle for diameter ℓ,

$$\alpha = \frac{\pi}{4}\ell\bar{\mathbf{A}}R(\Phi)\begin{bmatrix}1\\0\end{bmatrix}, \quad \beta = \frac{\pi}{4}\ell\bar{\mathbf{A}}R\left(\Phi+\frac{\pi}{2}\right)\begin{bmatrix}1\\0\end{bmatrix}, \tag{2.18}$$

and δ and Δ are the rates at which the local connection changes across the shape space, multiplied by the diameter of the gait,

$$\delta = \ell\frac{\partial \mathbf{A}_\alpha}{\partial r_\beta}, \quad \Delta = \ell\frac{\partial \mathbf{A}_\beta}{\partial r_\alpha}. \tag{2.19}$$

Inserting these approximations into the BCH series for the gait gives

$$\begin{aligned} g_\phi \approx \exp(&-\delta+\Delta+[\alpha,\beta]+ \\ &\frac{1}{2}[(\alpha+\beta),(-\delta+\Delta+[\alpha,\beta])]+\cdots, \end{aligned} \tag{2.20}$$

which we can then expand in terms of the local connection as

$$g_\phi \approx \exp\left(\underbrace{\overbrace{\iint_\phi \left(d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]\right)}^{-\delta+\Delta} + \overbrace{[\alpha,\beta]}^{}}_{\text{cBVI}} + \overbrace{\frac{\pi\ell}{8}\left[(\bar{\mathbf{A}}_\alpha + \bar{\mathbf{A}}_\beta), \iint_\phi \left(d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]\right)\right]}^{\text{third order effects}}\right). \tag{2.21}$$

This surface integral formulation opens the possibility of gait optimization via variational techniques. These specific optimization techniques are outside the scope of this paper; however, they assume that the cBVI is an accurate estimate of displacement (as in the optimized coordinates shown in the last row of Fig. 2.1). Our focus in this paper is the validity of this assumption, and quantifying residual errors due to the truncation of the BCH series. In particular, we express and bound the contribution of third order effects to displacement.

2.4. Third Order Bound

The third order effects in (2.21) depend on $(\alpha+\beta)$ and the cBVI. As in (2.18), the gait diameter ℓ and starting phase Φ together encode the initial configuration of the system. Within a fixed diameter gait, Φ is solely responsible for initial system configuration. In general, third order effects depend on the size and orientation of the local connection $\mathbf{A}$, the gait diameter ℓ, and starting phase Φ, as shown in Fig. 2.1.

2.4.1 Approximation

Third order effects (2.21) are a third-degree polynomial in ℓ; to produce a third order bound, we construct a similar polynomial approximation for the cBVI. This is done by computing a two-dimensional, second order Taylor series approximation for the total

Lie bracket, $D\mathbf{A}$, at the center of the gait:

$$D\mathbf{A} \approx \overline{D\mathbf{A}}(\delta\alpha_1, \delta\alpha_2) = T^2_{D\mathbf{A}}(\delta\alpha_1, \delta\alpha_2). \tag{2.22}$$

We then reparameterize (2.22) into polar coordinates, and integrate over the circular approximation for a gait,

$$\overline{\text{cBVI}}(\ell) = \int_0^{\ell/2} \int_0^{2\pi} \overline{D\mathbf{A}}(\rho, \theta)\rho \, d\rho \, d\theta, \tag{2.23}$$

producing a diameter-dependent estimate for the cBVI.

As third order effects (and associated estimates) depend on phase Φ and characteristic gait diameter ℓ, these parameters are used to compute a third order bound.

2.4.2 Third Order Bound

The preceding polynomial approximations assume a circular gait, and express both nominal and approximate displacements as a result. We use these approximations to construct a heuristic on the size of third order contributions for generic systems, and determine third order effects in the worst case.

The magnitude of third order effects can be made large by maximizing the possible size of its constituent components. This is first done with the triangle inequality on α and β, creating an upper bound on $(\alpha + \beta)$:

$$(\alpha + \beta) \leq (|\alpha| + |\beta|); \quad (|\alpha| + |\beta|) = \frac{\pi}{4}\ell|\bar{\mathbf{A}}| \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \tag{2.24}$$

Note that with the absolute values of α and β, phase Φ is no longer present. This implies that the bound captures the largest possible third order effects across all phases.

We also maximize the local Lie bracket, using the triangle inequality. This has the form

$$[X, Y] = \begin{bmatrix} X^y Y^\theta - Y^y X^\theta \\ Y^x X^\theta - X^x Y^\theta \\ 0 \end{bmatrix} \leq \begin{bmatrix} |X^y Y^\theta| + |Y^y X^\theta| \\ |Y^x X^\theta| + |X^x Y^\theta| \\ 0 \end{bmatrix}. \tag{2.25}$$

Combining the upper bounds, we have

$$\text{third order bound} = \frac{1}{2}[(|\alpha| + |\beta|), \overline{\text{cBVI}}(\ell)]_{\text{ub}}, \tag{2.26}$$

where $[\cdot, \cdot]_{\text{ub}}$ refers to the bound on the local Lie bracket.

An important note is that the local Lie bracket in (2.26) speaks to the direction of third order effects. For gaits with no net rotation, third order effects are orthogonal to the cBVI; this is demonstrated in Fig. 2.1 and Fig. 2.5. As a result, we may speak of third order effects in terms of the "error angle" they produce. Ground truth displacements lie on an arc with an equivalent angle, and a radius equivalent to the cBVI. Location on the arc is determined by the starting phase Φ of the particular gait.

2.4.3 Characteristic Length Bound

The third order bound is an increasing function of the characteristic diameter ℓ and the local connection $\mathbf{A}$; third order contributions are small if ℓ is small. The definition of "small" is relative, and is determined by the size of the local connection, which depends on the choice of body frame. For a given choice of coordinates, solving the inequality

$$[(|\alpha| + |\beta|), \overline{\text{cBVI}}(\ell)]_{\text{ub}} \leq P \cdot \overline{\text{cBVI}}(\ell) \tag{2.27}$$

for ℓ will constrain third order effects to a proportion P of the cBVI. Because all the quantities involved are polynomials in ℓ, (2.27) can be solved numerically.

2.5. Application of Bound

We now apply the techniques introduced in §2.4. for two systems, investigating particular gait families. We explore the direction and magnitude of third order effects in both original and minimum perturbation coordinates.

2.5.1 Systems

We investigate two example systems: the differential-drive car and the Purcell swimmer. Illustrations of each are in Fig. 2.4. Both reside in the plane and have two shape variables. For the car, the shape variables are the orientations of the wheels; for the swimmer, they are the relative orientations of every two links. A "shape" is a particular value for both shape variables; it defines the configuration of a system.

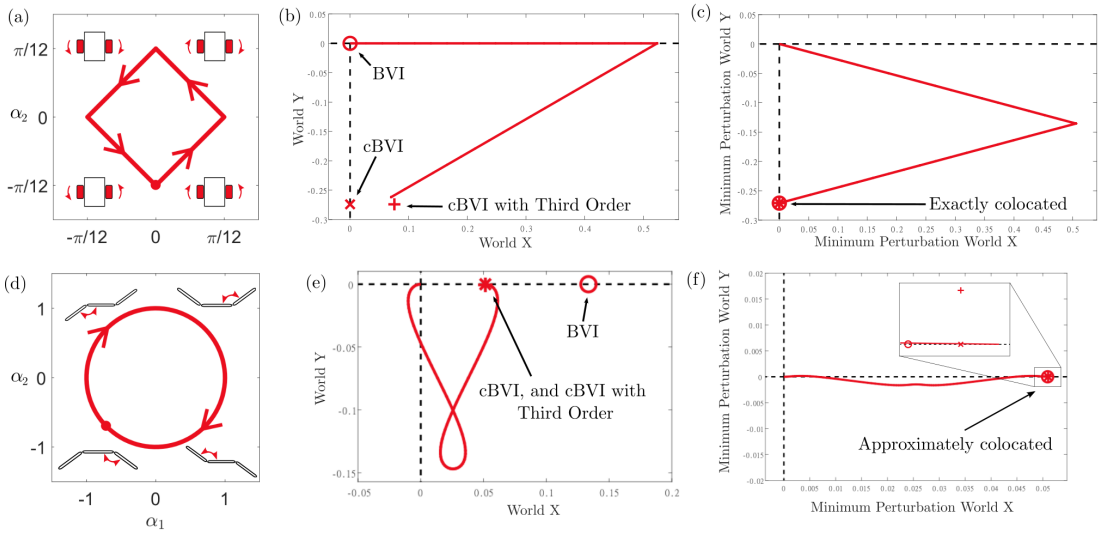

FIGURE 2.4: Characteristic gaits and resulting trajectories for each system. On the top row, (a-c), the diffdrive car moves in square gaits (a), with the orientation of each wheel as the shape. (b) Square diffdrive gaits result in a displacement trajectory. (c) In minimum perturbation coordinates, BCH estimates for displacement are exactly colocated. On the bottom row (d-f), the Purcell swimmer moves with circular gaits (d), with the relative orientation of each link as the shape. (e) Circular Purcell gaits result in a displacement trajectory. (f) In minimum perturbation coordinates, BCH estimates are only approximately colocated.

In general, the shape space of a system is all of the possible shapes it can make; gaits are closed loops within the shape space. In the case of the above systems, we can represent the shape space as a subset of $\mathbb{R}^2$, and draw closed loops on the plane to construct gaits of interest.

Fig. 2.4(a) and 2.4(d) show characteristic gaits investigated for each system. The diffdrive car has a square gait, as it executes discrete "move forward" and "turn" actions.

The Purcell swimmer has a circular gait, where it continuously accelerates each joint. Each gait results in a displacement through the world; this displacement (and how it varies for different gaits) is of principal interest.

2.5.2 Displacement and Effect of Additional Terms

Changes in system shape induce a displacement trajectory through position space, shown in Fig. 2.4(b) and 2.4(e).

As mentioned in §2.2., the ground truth is calculated exactly with a path integral of the local connection along the gait. Approximations of displacement are done with a surface integral of BCH terms inside the path; as more terms are included, the approximation becomes more accurate.

2.5.3 Minimum Perturbation Coordinates

Choice of body coordinate affects the trajectory that systems follow through position space. Body displacements may be computed in any body frame, so long as the frame is rigidly attached to the system. Minimum perturbation coordinates [14] are a choice of frame with this property.

Fig. 2.4(c) and 2.4(f) show how the use of minimum perturbation coordinates affects each system's trajectory. For the diffdrive car, displacement estimates are exact in minimum perturbation coordinates [12,14]: the BVI,[6] cBVI, and third order estimates are perfectly colocated. For the Purcell swimmer, the BVI, cBVI, and third order estimates are approximately colocated; in this case, they are 0.0031, 0.0013, and 0.0022 from ground truth, respectively.

It is important to note that Fig. 2.4 ignores the effect of starting phase on displacement, which does appear in higher order terms. This effect is addressed in both §2.5.4 and §2.5.5.

[6]The Body Velocity Integral (BVI) [14] is a first order estimate for displacement, making Fig. 2.4 capture first, second, and third order estimates.

2.5.4 Relative Third Order Contribution

To demonstrate the third order bound, we sample the gait families for both the diffdrive car and Purcell swimmer over gait amplitude and starting phase. In addition, we compute resulting ground truth, cBVI, and third order effects for each gait, in original and minimum perturbation coordinates.

Fig. 2.5 shows the results of amplitude and period sampling for the diffdrive car, in original coordinates. The sampling is omitted in minimum perturbation coordinates, as the BVI exactly captures displacement for this system [12, 14]. Fig. 2.1 shows the same sampling for the Purcell swimmer, in both original and minimum perturbation coordinates.

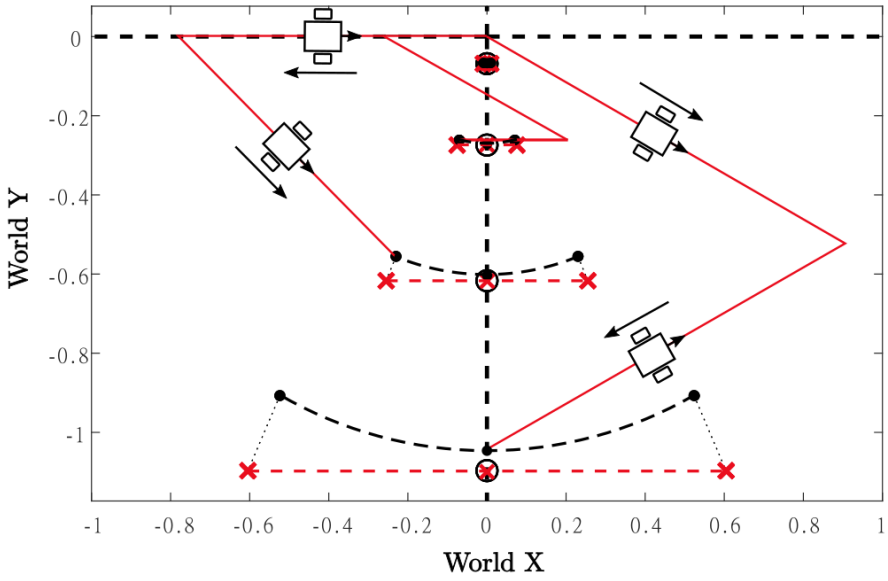

FIGURE 2.5: Amplitude and phase sampling for the diffdrive car, in original coordinates. Sample trajectories are included in red; resulting ground truths are in black. The cBVI captures changes in amplitude; third order effects (as red X's) capture changes in phase. Arc "error" angle increases with amplitude.

For both systems, it is clear that the magnitude of third order effects is much smaller than the cBVI, i.e., third order contributions to displacement are small (in optimal coordinates). The relative *size* of this contribution increases with gait amplitude. In cases where the bound is not acceptably small, it can be made so with the constraints on

amplitude.

2.5.5 Guarantees using Third Order Bounds

Using the length bound defined in §2.4.3, the magnitude of third order effects may be *absolutely* constrained to an arbitrary proportion of the cBVI. The effect of this bounding technique is shown in Fig. 2.6. The maximum "error angle" increases with characteristic gait diameter.

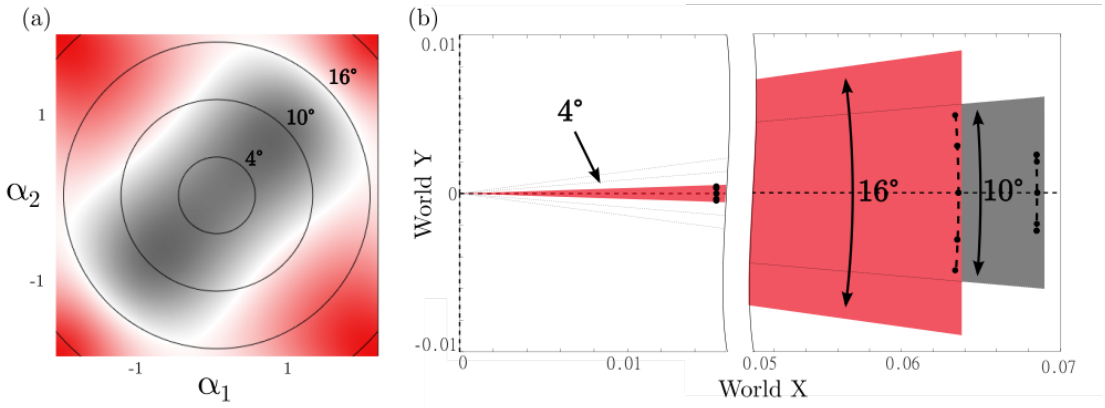

FIGURE 2.6: Error angle for the Purcell swimmer, in minimum perturbation coordinates. The X constraint curvature function is overlaid with gait contours; each gait diameter has an associated error angle (a). Error angle bounds the third order effects present in the ground truth (b). Ground truths (dashed lines) shown are of the same scale as in Fig. 2.1.

Third order effects may additionally be *relatively* constrained (within a given amplitude) with an intelligent selection of starting phase, Φ. As shown in (2.21), the local Lie bracket increases with $(\alpha + \beta)$; α and β are dependent on *phase*, as well as characteristic diameter. Certain choices of starting phase will result in relatively small third order effects. Fig. 2.1 and Fig. 2.5 corroborate this claim. Within a given amplitude, the sampled phases all have different third order contributions.

2.6. Conclusion

In this paper, we extend existing displacement approximations, and characterize third order effects of the BCH series in the context of locomoting systems. We identify that gait diameter, starting phase, and coordinate choice influence third order contributions, and demonstrate the use of these quantities as tools to manage errors introduced by the cBVI.

Future work will explore third order effects in the context of gaits with net rotation; these will act non-orthogonally to the cBVI, and require further analysis. In addition, we will expand scope to include fourth order terms, which capture additional, previously ignored displacement effects.

Acknowledgment

This work was supported in part by the National Science Foundation under grants 1653220 and 1826446.

Minimum Perturbation Coordinates on the 3D Rotation Group

Capprin Bass and Ross L. Hatton

Abstract

When considering mobile robot systems, displacement estimates like the constraint curvature function are useful for intuition of dynamics and optimization of control policies. These estimates are valid only for small shape changes, unless particular body coordinates (minimum perturbation coordinates) are used to compute displacements. In this paper, we extend prior work to compute minimum perturbation coordinates on the 3D rotation group, $SO(3)$. We improve the derivation of the objective function and introduce a more correct method for reconstruction of dynamics after coordinate change. We demonstrate these methods on an example system, and show that minimum perturbation coordinates are more effective on $SO(3)$ than previously thought.

Unpublished Manuscript

3.1. Introduction

A key modern problem in robotics is the construction and control of mobile locomotors. These are systems that move through the world using highly dynamic interactions with their environment; examples include legged locomotors [17–19], snakes [20, 21], and fluid swimmers [2, 22, 23]. Such systems have a broad set of dynamic effects, including contributions from inertia, coriolis terms [2], and storage or dissipation of energy [24, 25].

Classical mechanics provides the tools to model dynamic systems, in the sense that kinematics, forces, and torques may be mapped through the robot. However, these tools provide little intuition on a high level about which control actions are effective for specific locomotion objectives. More specifically, the relationship between internal changes in shape and resulting trajectory through the world is unclear.

The geometric mechanics community has defined tools that describe system dynamics at this high level. In particular, the Constraint Curvature Function (CCF) [12] provides intuition about which internal motions result in specific world displacements. This was not possible until the importance of the choice of body frame coordinates was well understood, and minimum perturbation coordinates were derived for planar systems [14].

By choosing minimum perturbation coordinates as the reference frame for mobile systems, we can use the CCF to visualize which parts of robot shape spaces produce desired motion through the world [12]. As natural extensions, the CCF has been used as an objective function to produce optimal control policies for dynamic systems [1, 26, 27] and do online dynamics estimation and control for systems with unknown geometry [28]. Minimum perturbation coordinates are also a natural choice when doing learning-based control, as they ensure that the learned policy is responsible for system performance, not the particular choice of body coordinate.

Previous work has derived minimum perturbation coordinates for planar systems

[14] as well as systems on the 3D rotation group, $SO(3)$ [29]. This paper addresses gaps in the latter work, specifically correcting, completing, and simplifying the derivation of minimum perturbation coordinates on $SO(3)$ in order to make the content more approachable for others. We construct an improved expression of the objective function, describe how optimization is performed, and explain how to reconstruct the local connection using the time derivative of the exponential map.

The remainder of this paper is organized as follows. In §3.2., we describe the relevant model and detail previous contributions towards minimum perturbation coordinates in $SE(2)$ and $SO(3)$. In §3.3., we express our contributions towards the derivation and computation of minimum perturbation coordinates on $SO(3)$. In §3.4., we demonstrate coordinate optimization results for a reaction wheel satellite, the same system demonstrated in [29]. In §3.5., we make concluding remarks and discuss future work.

3.2. Background

3.2.1 Model

We model mobile systems respecting their configuration space Q, partitioned into a shape space B and position space G. The shape of a system is defined by the configuration of its internal degrees of freedom. The position is defined by the coordinates of some reference frame attached to the system, and reflects which parameters are important to describe motion of the system. As an example, planar systems are often represented with $[x, y, \theta]^T \in G$, implying a special Euclidean position space: $G = SE(2)$. In this paper, the position space is the group of 3D rotations: $G = SO(3)$.

The distinction between shape and position space is necessary to develop a notion of how internal changes in shape produce displacement in the world. A motility map expresses this relationship, acting as a local connection between respective tangent spaces.

For $r \in B$, this is:

$$\overset{\circ}{g} = \mathbf{A}(r)\dot{r}, \tag{3.1}$$

where $\overset{\circ}{g} \in \mathfrak{g}$ is a body velocity, residing in the Lie algebra of the position space. The local connection may be constructed for kinematic systems using the nullspace of system constraints [2, 22], or for dynamic systems, given specific properties [15].

By considering each row of the local connection separately, we can construct vector fields over the shape space that map between shape velocities and resulting body velocities in the position space [30]. As an extension, motion plans in the shape space result in well-defined trajectories through the world; internal motion is mapped *forward* through the system dynamics. However, the reverse operation is more challenging. The local connection evaluated at a given shape does not necessarily have a unique inverse, and integration of its pseudoinverse over time does not respect the need to stay within joint limits. Additional structure is necessary to discuss closed gait geometry while analyzing the gaits that produce specific world displacements. This information is provided by the Constraint Curvature Function.

3.2.2 Constraint Curvature Function

To intuit which regions of the shape space are valuable for specific displacements, we extend the local connection to measure accrued differences in displacement due to motion (*nonconservativity*) as well as compound effects from motion in different dimensions (*noncommutativity*) [12]. Expressed mathematically, this is

$$g_\phi = \iint_\phi \overbrace{d\mathbf{A}}^{\text{nonconservativity}} + \underbrace{[\mathbf{A}_1, \mathbf{A}_2] + \text{higher order terms}}_{\text{noncommutativity}} \; d\phi, \tag{3.2}$$

where ϕ is a closed[1] gait cycle, g_ϕ is the world displacement induced by the cycle, and $[\cdot,\cdot]$ is a local Lie bracket of the columns of the local connection. The integrand in (3.2) is the Constraint Curvature Function (CCF) [12, 29], which may be evaluated over the shape space to indicate which regions produce desired displacement through the world.

As mentioned in [12,29,31], the CCF is a truncation of the Baker-Campbell-Hasudorff series, and ignores higher order displacement effects. These effects may be ignored for small motions; however, as gait size increases, the magnitude of higher order terms results in significant accrued error [31]. This is addressed by using a particular choice of body coordinates, which accounts for error by minimizing contributions to nonconservative effects.

3.2.3 Minimum Perturbation Coordinates on $SE(2)$

The general notion of choosing "minimum perturbation" coordinates is based on the relative size of the local connection with respect to different body frames. Intuitively, we pick a body frame such that contributions towards the integrand of (3.2) are as small as possible; this corresponds to minimizing $\mathbf{A}$ with respect to constrained choices of body frame. As described in [12, 14, 29] we can pick any frame that acts like a rigid body attached to the system.

Specifically, rotating the body frame changes only the curl-free component of the local connection vector field, while preserving the divergence-free component; the structure of $SE(2)$ causes rotations to commute. Referencing (3.2), we minimize noncommutative effects while preserving absolute nonconservativity (the $d\mathbf{A}$ term) [14]. The objective function describing minimum rotation in $SE(2)$ is

$$D = \iint_\Omega || \underbrace{\mathbf{A}^\theta(r) + \nabla_r b_\theta(r)}_{\mathbf{A}^\theta_{\text{opt}}(r)} ||^2 d\Omega, \tag{3.3}$$

[1]The closed-loop constraint is a result of replacing a line integral (in the case of the local connection) with an area integral. Formally, this is an expression of Green's theorem (for 2D shape spaces) or Stokes' theorem. Intuitively, this is because nonconservativity, in a similar sense as voltage, is a potential expressed as a difference between two points.

where we minimize D with respect to a shape-varying coordinate transform $b(r)$ in exponential coordinates within a region Ω of the shape space [14].

On $SE(2)$, the choice of body frame $\exp(b)$ is optimized numerically using a discrete Hodge-Hemholtz decomposition, which respects the curl- and divergence-free components of the local connection [14]. On $SO(3)$, however, noncommutativity between rotations on different axes necessitates a different problem formulation.

3.2.4 Minimum Perturbation Coordinates on $SO(3)$

The motivation behind minimum perturbation coordinates on $SO(3)$ is the same as on $SE(2)$: we look for a minimal local connection that preserves system behavior. However, rotations in 3D are noncommutative, meaning that series' of rotations along different axes have different results depending on order.[2] As a result, the expression for the optimized local connection on $SO(3)$ contains terms that capture noncommutativity in rotation:

$$\mathbf{A}_{\text{opt}}(r) = \mathbf{A}(r) - \underbrace{b(r)\mathbf{A}(r)}_{\text{noncommutativity}} - \nabla_r b(r), \tag{3.4}$$

where $b(r)$ is still a shape-varying coordinate transform in exponential coordinates [29].

A second consequence of noncommutativity on $SO(3)$ is that rotations are coupled; that is, rotations that act along two axes induce additional rotation along the third axis. Coordinate optimization is performed simultaneously along each dimension to address coupled rotation:

$$D = \iint_{\Omega} ||\mathbf{A}^x_{\text{opt}}(r)||^2 + ||\mathbf{A}^y_{\text{opt}}(r)||^2 + ||\mathbf{A}^z_{\text{opt}}(r)||^2 d\Omega, \tag{3.5}$$

where Ω is still a region of the shape space, and $\{x, y, z\}$ are the axes of rotation. The transform $b(r)$ is optimized numerically by finding weights on linear basis functions defined over a discretization of the shape space [29].

[2] As an example, imagine a 90° rotation of any coordinate frame about its x axis, followed by another 90° rotation about its y axis. The result is quite different than the same rotations applied in the reverse order.

The details on the derivation and optimization of the objective function presented in previous work [29] are vague, and in some cases, not correct. In this paper, we reformulate minimum perturbation coordinates on $SO(3)$, fully describing the objective function and process of optimization. In addition, we contribute a more complete and correct expression of the optimized local connection $\mathbf{A}_{\text{opt}}(r)$ using the time derivative of the exponential map.

3.3. Minimum Perturbation Coordinates on the 3D Rotation Group

3.3.1 Derivation

Minimum perturbation coordinates are given by a transform $\beta(r) \in SO(3)$ that changes over the shape space. We express this transform using exponential coordinates, represented by a body velocity $b(r) \in \mathfrak{so}(3)$ on the Lie algebra of $SO(3)$. The relationship between velocities expressed in original and optimal coordinates is given instantaneously by

$$\mathring{g}_{\text{opt}} = \text{Ad}_{\beta}^{-1}\big(\underbrace{\mathring{g}_{\text{orig}}}_{-\mathbf{A}(r)\dot{r}} + \nabla_r b(r)\dot{r}\big), \tag{3.6}$$

where velocity in the original frame $\mathring{g}_{\text{orig}}$ is transferred to the optimal frame, accounting for contributions from changes in $b(r)$ in the new frame.

On $SO(3)$, the adjoint mapping is equivalent to a corresponding spatial rotation. In our case, we can represent the adjoint Ad_{β}^{-1} as $\exp(-b)$. Applying these identities to (3.6), we construct a local connection in the optimal frame:

$$\mathring{g}_{opt} = \underbrace{\exp(-b)\big(-\mathbf{A}(r) + \nabla_r b(r)\big)}_{-\mathbf{A}_{\text{opt}}(r)}\dot{r}. \tag{3.7}$$

The local connection described in (3.7) is nonlinear in b, making minimization with respect to rigid body constraints challenging. Instead, we perform a linearization of the exponential map, which will allow us to optimize β over a discretization of the shape

space. The first order approximation for exp is

$$\exp(X) = I + \rho(X), \tag{3.8}$$

where ρ is the general linear matrix representation for X. Applied to (3.7) and dropping both higher order terms and function notation, we have

$$\mathring{g}_{\text{opt}} \approx \underbrace{\big(-\mathbf{A} + [-\mathbf{A}, b] + \nabla_r b \big)}_{-\mathbf{A}_{\text{opt}}(r)} \dot{r}. \tag{3.9}$$

Note that noncommutativity between rotations is captured by the Lie bracket term. When expressed for each dimension, (3.9) is

$$\mathbf{A}^x_{\text{opt}} = \mathbf{A}^x + b^y \mathbf{A}^z - b^z \mathbf{A}^y - \nabla_r b^x, \tag{3.10}$$

$$\mathbf{A}^y_{\text{opt}} = \mathbf{A}^y + b^z \mathbf{A}^x - b^x \mathbf{A}^z - \nabla_r b^y, \tag{3.11}$$

$$\mathbf{A}^z_{\text{opt}} = \mathbf{A}^z + b^x \mathbf{A}^y - b^y \mathbf{A}^x - \nabla_r b^z. \tag{3.12}$$

The above result is similar to (but not the same as) the expression for the optimized local connection provided by [29]; sign changes on the gradient terms indicate inconsistencies in the previous work. We believe that this is the more correct expression.

The objective function is the same as in (3.5), with the goal of simultaneously minimizing the magnitude of all components of the local connection. We now describe the computational approach to minimization of the objective function.

3.3.2 Computation

As in [29], b is expressed as a weighted sum of linear basis functions over a discrete mesh with n vertices. That is, $b^x = \sum_i^n f_i^x e_i^x$, $b^y = \sum_i^n f_i^y e_i^y$, and $b^z = \sum_i^n f_i^z e_i^z$, where f_i is a linear function of unit height at point i, and e_i is the associated weight. Optimization consists of solving for weights e_i^j such that the resulting b minimizes (3.5).

We minimize (3.5) by finding b such that

$$\frac{\partial D}{\partial e^x} = \frac{\partial D}{\partial e^y} = \frac{\partial D}{\partial e^z} = 0. \tag{3.13}$$

The linear finite element approach allows (3.13) to be expressed as a linear equation, and weights found as a solution of said equation. This expression is

$$\int_\Omega \begin{bmatrix} \nabla f_i \cdot \nabla b^x + f_i b^x \left(\mathbf{A}^z \cdot \mathbf{A}^z + \mathbf{A}^y \cdot \mathbf{A}^y\right) & \nabla f_i \cdot b^y \mathbf{A}^z - f_i \mathbf{A}^z \cdot \nabla b^y & -\nabla f_i \cdot b^z \mathbf{A}^y + f_i \mathbf{A}^y \cdot \nabla b^z \\ \nabla f_i \cdot b^x \mathbf{A}^z + f_i \mathbf{A}^z \cdot \nabla b^x & \nabla f_i \cdot \nabla b^y + f_i b^y \left(\mathbf{A}^z \cdot \mathbf{A}^z - \mathbf{A}^x \cdot \mathbf{A}^x\right) & -\nabla f_i \cdot b^z \mathbf{A}^x - f_i \mathbf{A}^x \cdot \nabla b^z \\ \nabla f_i \cdot b^x \mathbf{A}^y - f_i \mathbf{A}^y \cdot \nabla b^x & -\nabla f_i \cdot b^y \mathbf{A}^x + f_i \mathbf{A}^x \cdot \nabla b^y & \nabla f_i \cdot \nabla b^z + f_i b^z \left(\mathbf{A}^y \cdot \mathbf{A}^y + \mathbf{A}^x \cdot \mathbf{A}^x\right) \end{bmatrix} = \int_\Omega \begin{bmatrix} \nabla f_i \cdot \mathbf{A}^x \\ \nabla f_i \cdot \mathbf{A}^y \\ \nabla f_i \cdot \mathbf{A}^z \end{bmatrix}, \quad (3.14)$$

and is equivalent to that given in [29]. Gradients are computed numerically over the grid, and integration is performed with quadrature (respecting the fact that functions over the grid are all linear).

Note the following. Each row of (3.14) corresponds to a single x, y, or z dimension. Each column corresponds to contributions from x, y, and z directions of the tranform b. Because there are n points in our mesh and b^j contains n weighted basis functions, each entry is a $n \times n$ block of values.

Because (3.14) depends linearly on weights e_i^j, we can trivially express the optimal transform $b(r)$ as the solution to said linear equation.

3.3.3 Reconstructing the Local Connection

Given the optimal transform $b(r)$ in exponential coordinates, the only remaining task is to construct the local connection $\mathbf{A}_{\text{opt}}$ in the optimized frame. Naïvely, this is done with the instantaneous expression for the local connection given in (3.7) or its linearization given in (3.9). We believe [29] uses the former definition. The more correct approach takes advantage of the full expression, which accounts for changes in b along its exponential flow. We explain this process here.

We can interpret (3.7) as the velocity of the optimal frame g_{opt} at the end of the

flow $\exp(b)$. The complete expression is

$$\overset{\circ}{g}_{\text{opt}} = \text{Ad}_{\beta}^{-1}\left(\overset{\circ}{g}_{\text{orig}}\right) + \overset{\circ}{\exp}(b), \tag{3.15}$$

where $\beta = \exp(b)$ is the transform generated by b. The contribution from b is written as

$$\overset{\circ}{\exp}(b) = \exp(-b)\left(\frac{d}{dt}\exp(b)\right), \tag{3.16}$$

which accounts for how changes in b propogate into the optimal frame. Because b changes along its exponential flow, we integrate change along the flow to find $\overset{\circ}{\exp}(b)$:

$$\overset{\circ}{\exp}(b) = \int_0^1 \text{Ad}_{\beta^{-1}(\tau)\beta(1)}^{-1}\left(\frac{d}{dt}b\right) d\tau. \tag{3.17}$$

We express $\overset{\circ}{g}_{\text{opt}}$ using this contribution from b:

$$\overset{\circ}{g}_{\text{opt}} = \text{Ad}_{\beta}^{-1}\underbrace{\left(\overset{\circ}{g}_{\text{orig}}\right)}_{-\mathbf{A}_{\text{orig}}(r)\dot{r}} + \int_0^1 \text{Ad}_{\beta^{-1}(\tau)\beta(1)}^{-1}\underbrace{\left(\frac{d}{dt}b\right)}_{\nabla_r b(r)\dot{r}} d\tau. \tag{3.18}$$

If we do not account for change in b along the flow, (3.18) simplifies to the expression given in (3.7). This is a reasonable choice when optimizing for b, as we only use local information for that process. However, given $b(r)$ after optimization, it is important to use this complete expression for $\overset{\circ}{g}_{\text{opt}}$ to capture all the effects of being in the optimal frame. We demonstrate the impact of this approach in our results.

3.4. Results and Analysis

Here, we compute minimum perturbation coordinates for an example system on the 3D rotation group. We compare the resulting displacement maps to those in original coordinates and prior work.

3.4.1 Reaction Wheel Satellite

The example system for this paper is the same reaction wheel satellite used in previous work [29], shown in Fig. 3.1. The body has wheels at fixed distance from the

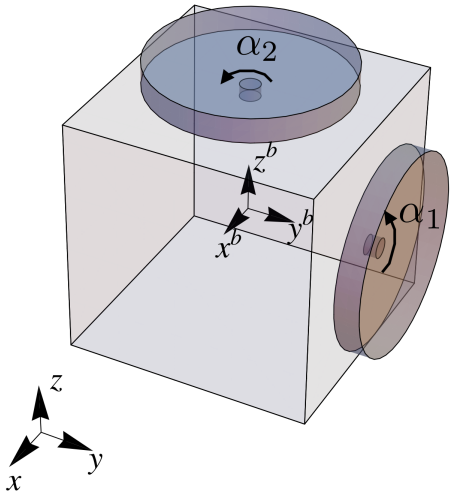

FIGURE 3.1: Reaction wheel satellite, duplicated from [29]. Wheels centered on the y and z body axes induce body rotations about the center of the system.

body frame, acting about the y and z axes. Rotations about each axis are parametrized by α_1 and α_2, respectively, and induce a counterrotation of the body about the same axis. As a result, the shape space of this reaction wheel satellite is 2D, and the position space of this system is $SO(3)$. We compute minimum perturbation coordinates for this system, and compare displacement maps in original and optimal coordinates.

3.4.2 Optimal Displacement Maps

Fig. 3.2 shows the local connection for the reaction wheel satellite in original coordinates, as well as coordinate optimization results from [29] and this paper. For the same input behavior and shape space, coordinate optimization results are qualitatively similar between this work and the previous work. However, the optimized local connection vector fields are clearly different. We credit this to the methods for reconstructing the local connection provided in §3.3.3, which reveal different structure than that in previous work.

The differences between optimization results are even clearer when we compare the CCF and displacement from an example gait. Fig. 3.3 compares the CCF from [29] to that from this paper, including the estimated displacement from an example gait. The general shape of the CCF is again similar for both optimization results, indicating similar

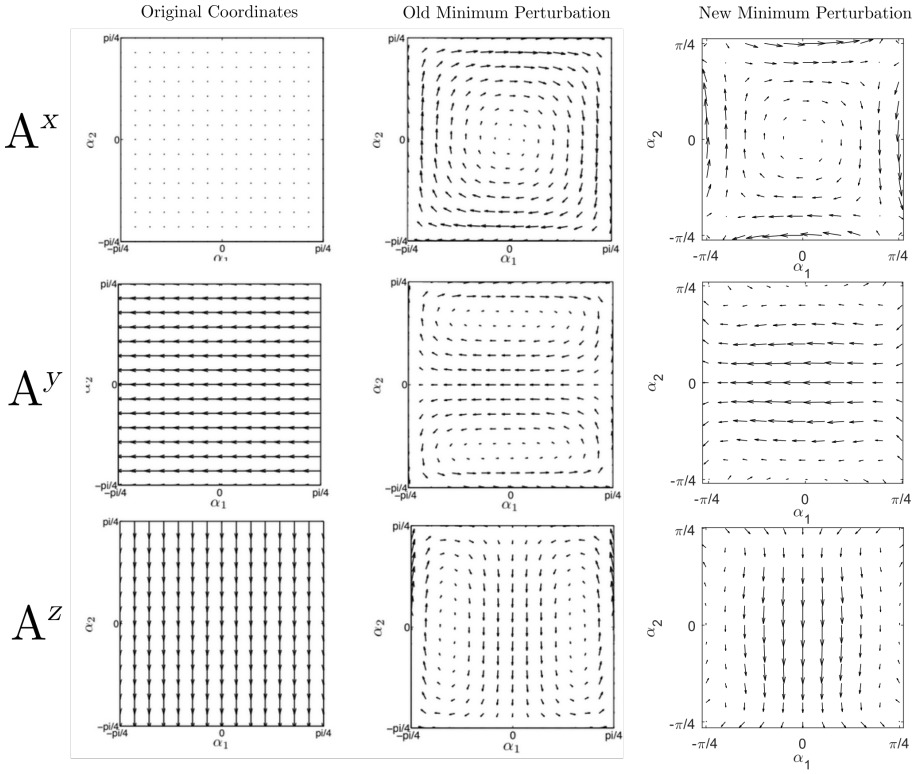

FIGURE 3.2: Local connection vector fields in original and optimized coordinates; the first two columns are duplicated from [29] for comparison. The behavior after optimization is generally the same for current and prior work. Differences in reconstruction of the local connection provided in §3.3.3 reveal slightly different behavior of the local connection after optimization.

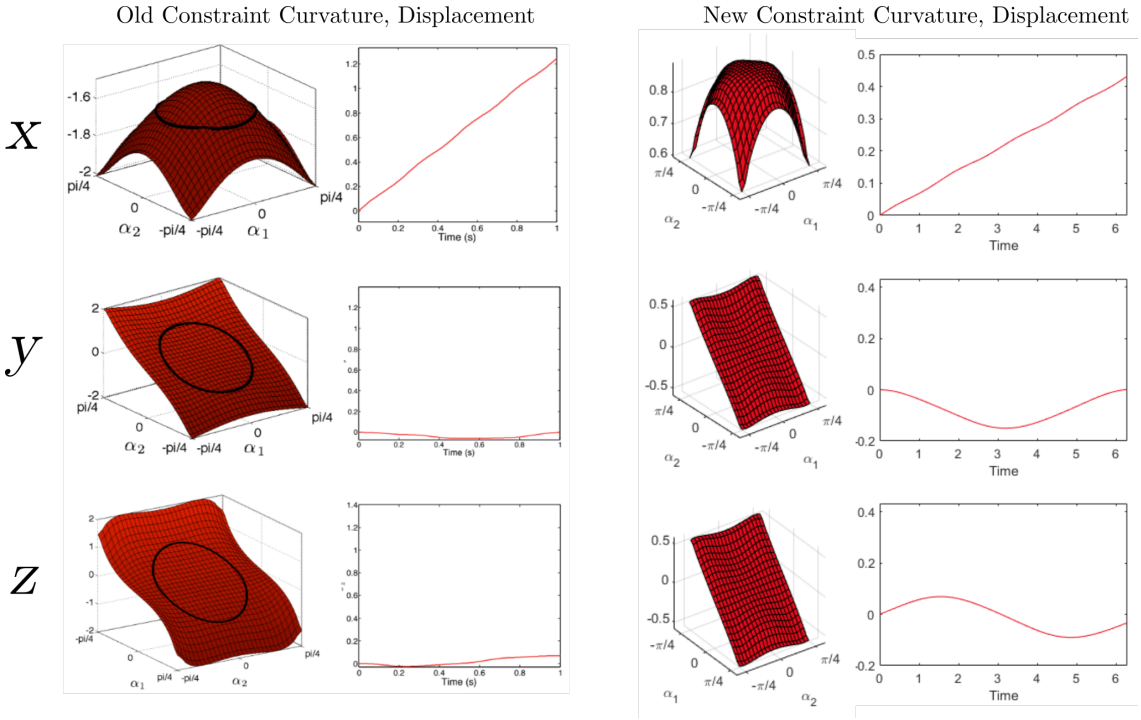

FIGURE 3.3: CCF and predicted displacement for an example gait; old figures duplicated from [29]. With the improvements to reconstructing the local connection from §3.3.3, magnitudes of both the CCF and resulting displacements are much smaller than in prior work. For the new figures, the gait is omitted, and time is in units of radians.

behavior to that found in [29]. However, the magnitude of the CCF produced in this paper is significantly smaller for all dimensions, indicating that coordinate optimization has better results than published previously. The improved coordinate optimization results are corroborated by the predicted displacement across the same example gait; across the board, the magnitude of displacement in each dimension is either very similar to or smaller than previous work.

In general, the results from this paper corroborate those of [29]. Our improved method of reconstructing the local connection slightly changes the local connection vector fields, and shows that the magnitude(s) of displacement after coordinate optimization are smaller than demonstrated previously.

3.5. Conclusion

The primary purpose of this work is to clarify the expression of coordinate optimization on $SO(3)$. We augment the coordinate optimization methods introduced in [29] with further explanation, including motivations that span many prior papers. We also provide a clearer derivation of minimum perturbation coordinates, including an improved method of reconstructing the local connection after optimization. We implement our methods on an example system, and demonstrate that minimum perturbation coordinates perform better than shown in previous work.

There is a significant technical barrier to the methods explored here; we hope that this document helps to make the content introduced in [29] more approachable. Future work could investigate nonlinear optimization methods; these include expansion of the exponential map to higher order terms, using the full expression for the optimized local connection, or exploring the Cayley transform to linearize $SO(3)$ before optimization. These approaches would likely reveal further improvements in performance, similar to those shown here.

Geometric Nonlinear Dimensionality Reduction for Robot Shape Spaces

Capprin Bass and Ross L. Hatton

Abstract

For robot systems with many degrees of freedom, much of the control space is often uninteresting or not useful. In addition, data in this space (scalar, vector, or higher-order fields) is challenging to visualize; as a result, it is difficult to explain or defend control policies. Building on our own previous work, we construct a method for performing dimensionality reduction on kinematic robot shape spaces. The resulting 2D space has well-defined mappings for both intrinsic structures and extrinsic functions, and so is a viable candidate for visualization of robot capabilities. Construction of this space is demonstrated for the family of discrete swimmers to act as example systems, and the results and implications are discussed.

Unpublished Manuscript

4.1. Introduction

Many robot systems today are designed to imitate biological organisms. On the macroscopic scale, these are dynamic locomotors, such as Cassie [17], Mini Cheetah [18], and ATLAS [19]. On smaller scales, systems imitate snakes [32] and swimmers [22]. A common feature of biomimetic systems is that they often have many (4+) degrees-of-freedom (DoF). This high-dimensional shape space is impossible to visualize at full complexity, and functions over this space (scalar fields, vector fields, etc) inherit the same problem.

Many approaches may be used to visualize high dimensional data [33]; however, these are general, and fail to capture the specific features and behavior of mobile systems. To remedy this, dimensionality reduction and visualization techniques must take into account system geometry and mechanics. Doing so will effectively capture the most important regions and behavior of the control space.

Prior work from the geometric mechanics community [12, 34] defines clear measures for "good" regions of robot shape spaces for specific locomotion objectives, e.g. displacement in particular local directions. These measures have been used in our previous motion planning work [1] to compute optimal policies in polynomial time, under constraints for direction traveled and motion cost.

In this paper, we present a method to reduce the dimensionality of robot shape spaces to 2D by using prior geometric tools (including variational gait optimization) to identify an optimal choice of subsurface in the shape space. Surfaces may be generated with respect to any objective function. An example is shown in Fig. 4.1. This choice of subsurface reduces the shape space of the robot, while preserving its most useful capabilities (forward displacement). Functions that are defined for the entire robot are still defined in this reduced space; conversely, structures (motion plans) in the reduced space

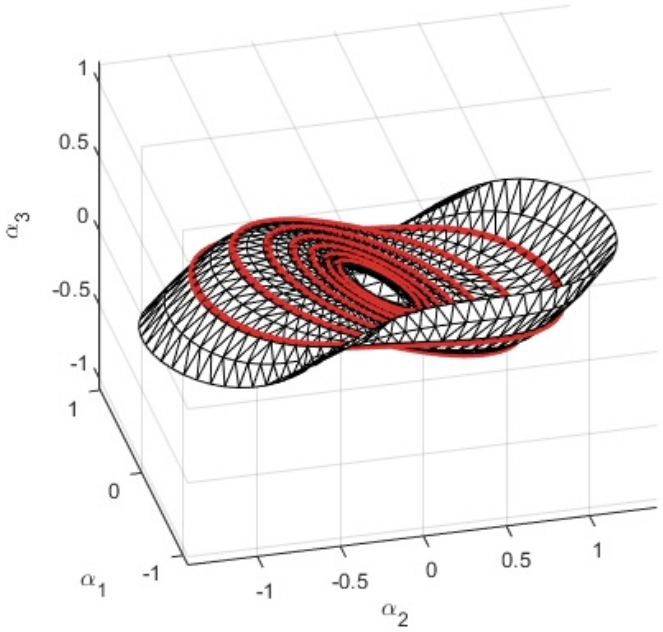

FIGURE 4.1: Reduced shape space, constructed for a 4-link drag-dominated swimmer. The reduced space is a 2D surface, constructed with a union of efficiency-optimal gaits (shown in red) for forward displacement. Motion plans on the surface are valid in the full, 3D shape space.

may be mapped up to the entire robot. As a result, this choice of reduced space is still valid for consideration of robot motion, and quantities in the complete robot shape space may be *visualized* in the reduced space.

The geometric dimensionality reduction approach builds on our motion planning work, which specifically produces single cyclic motion plans (gaits) [1]. Current work [27] extends variational gait optimization, identifying the family of gaits that are optimal with respect to a specified objective function. The resulting gaits are intrinsically one-dimensional (parametrized by one variable, t). The union of this gait family, Φ, over increasing gait diameter ℓ, produces an intrinsically two-dimensional surface immersed in the shape space,

$$S = \cup_{\ell,t}\Phi; \quad S \subset B, \tag{4.1}$$

where S is the two-dimensional space, and B is the shape space. Fig. 4.1 demonstrates

how a gait family constructs the immersed surface. Smoothness of the underlying functions implies that this space is continuously differentiable; so, there exists an injective mapping between the respective tangent spaces of the surface and the shape space. This property allows structures to be mapped from the surface to the shape space, or functions projected from the shape space to the surface.

This paper is organized as follows. In §4.2., we review the underlying model and relevant prior work. In §4.3., we construct the reduced shape space, and discuss its properties. In addition, we provide details for implementation in-code. In §4.4., we demonstrate the dimensionality reduction approach on a set of example systems: viscous and free-floating swimmers with many DoF. In §4.5., we discuss the important consequences of this approach, and comment on future work.

4.2. Background

4.2.1 N-Link Swimmers

Throughout this paper, we reference the 3-link and 4-link "swimmers" as examples of robot systems. As shown in Fig. 4.2, these systems are parametrized by 2 and 3 shape variables respectively, and reside in the plane. The motion of the N-link swimmers depends on the surrounding medium; here, we explore swimmers that are either *drag-dominated* (in a viscous fluid) or *isolated* (in free fall, or residing in a vacuum).

4.2.2 Geometric Mechanical Model

We model our systems by partitioning the configuration space Q into a position space G and shape space B.

The position space reflects the unique positions of the system in the world, often in terms of translation and orientation. In this paper, we primarily use $G = SE(2)$, with elements $[x, y, \theta]^T \in G$; however, these methods also generalize to other possible position

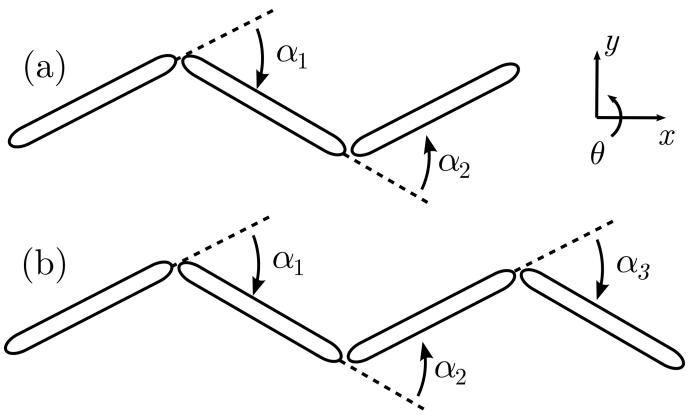

FIGURE 4.2: Illustration of planar N-link swimmers. (a) The 3-link swimmer is parametrized by two shape variables. (b) The 4-link swimmer is parametrized by three shape variables.

spaces [29, 32].

The shape space B captures the internal DoF of the system, and is often of the same dimension. One natural parametrization of the shape space is the actuators on a system; here, this is the relative orientations of "links" in the swimmer's body. Multiple shape variables may be combined to produce a smaller set of *shape modes*, which still parametrize system behavior with less complexity.

For the analysis in this paper, we assume that systems behave kinematically;[1] as a result, we can use the the *local connection* to map internal robot motion (in the shape space) to displacement through the world. The details of constructing the local connection are left to [35] and [2]. We map between elements $\alpha \in B$ and $g \in G$:

$$\mathring{g} = \mathbf{A}(\alpha)\dot{\alpha}, \tag{4.2}$$

where $\mathbf{A}$ is the local connection, and $\mathring{g}$ is the body velocity[2] resulting from $\dot{\alpha}$. As in [35] and [2], (4.2) may be integrated over motion plans (trajectories in the shape space) to compute displacement through the world resulting from specific changes in internal robot

[1]The kinematic domain applies to a number of systems, including drag-dominated [35] and inertial [2] systems. These can be swimming systems, snakes, or other low-speed crawlers.

[2]Body velocities are on the Lie algebra of the position space G.

shape.

Because mobile systems often have joint limits, it is convenient to express motion plans as *gaits* in the shape space. While displacement can still be computed using a line integral, the *total Lie bracket* is a more useful tool, as it can be used to define a gradient of displacement with respect to the gait parametrization. The expression for diplacement using the total Lie bracket is

$$g_\phi = \exp\left(\iint_\phi \overbrace{d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]}^{\text{total Lie bracket}} d\phi\right), \tag{4.3}$$

where ϕ is a gait parametrization, and g_ϕ is the displacement resulting from the gait. The total Lie bracket, when sampled across the shape space, is referred to as the *Constraint Curvature Function* (CCF), and its gradient is used alongside variational techniques to optimize gait shape for displacement [1].

4.2.3 Bilinear Form

The Jacobian $J : T_S \to T_B$ is a linear map between the tangent spaces of a shape space B and a subspace S. It can perform several key tasks for our purposes. In a *pushforward*, J maps intrinsically-defined structures (shapes, velocities, gaits, etc) from S to B. The pushforward guarantees that motion plans in a reduced shape space are valid in the full shape space. A *pullback* uses J to project extrinsically-defined functions from B onto S. The pullback allows us to define objective functions for the entire robot (such as the CCF), and express them in the reduced space.

In this paper, we pullback the *bilinear form* of the CCF using the Jacobian. The

CCF Ω^B, defined in B, is represented using a skew-symmetric matrix:

$$\Omega(v_1^B, v_2^B) = (v_1^B)^T \underbrace{\begin{bmatrix} 0 & \Omega_{12} & \Omega_{13} & \cdots \\ -\Omega_{21} & 0 & \Omega_{23} & \cdots \\ -\Omega_{13} & -\Omega_{23} & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}}_{\Omega^B} v_2^B, \tag{4.4}$$

where $v_{i,j}^B$ are vectors of interest for evaluating the CCF. The Jacobian J is used to pushforward velocities, or equivalently, pullback the extrinsically-defined CCF:

$$\Omega(v_1^S, v_2^S) = \overbrace{(v_1^B)^T}^{} \; J^T \underbrace{\begin{bmatrix} 0 & \Omega_{12} & \Omega_{13} & \cdots \\ -\Omega_{12} & 0 & \Omega_{23} & \cdots \\ -\Omega_{13} & -\Omega_{23} & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}}_{\Omega^G} \overbrace{J \; v_2^S}^{v_2^B}. \tag{4.5}$$

4.3. Dimensionality Reduction

Here, we describe our dimensionality reduction approach. The objective is to construct a reduced (2D) robot shape space for motion planning that preserves robot capabilities. Structures (shapes, gaits, or other plans) in this domain must map to the full shape space; conversely, it must be possible to express functions defined over the full shape space (like the CCF) in the reduced shape space.

We first describe the underlying mathematics of the reduced space, which give rise to our desired properties. Then, we discuss the optimal surface and associated mappings. Finally, we cover details of implementation in-code.

4.3.1 Mathematical Underpinnings

Gaits are parametric functions, mapping a single variable into the shape space. As an example, for an execution time T, parameter $t \in [0, T)$, and $\alpha \in B$, we define a gait ϕ:

$$\phi : [0, T) \to B; \quad \phi(t) = \alpha. \tag{4.6}$$

The parametric gait definition implies that gaits are intrinsically one-dimensional, even if the (extrinsic) shape space B has many dimensions.

Another property is that an increase in position space displacement requires an increase in gait size, corresponding to the area enclosed by the gait [31]. We refer to the size of a gait by its average diameter, ℓ.

If we define a gait family Φ as a set of gaits at all resulting displacements (equivalently, gait diameters):

$$\phi(\ell; t) \in \Phi, \tag{4.7}$$

then the *ordered union*[3] of Φ over ℓ and t is a set defined on B. That is,

$$S = \cup_{\ell,t}\Phi; \quad S \subset B. \tag{4.8}$$

This process, illustrated in Fig. 4.3, defines the set S. Any element of S has a neighborhood defined by ℓ (neighbors across gaits) and t (neighbors within gaits). The functions that parametrize S are all continuous and differentiable,[4] implying that its gradient with respect to the basis directions in B is always defined.

So, the union of gaits across displacement is an intrinsically 2D surface, immersed

[3]Here, an ordered union refers to a standard union of sets, but imposes an order relation based on ℓ and t. In effect, points are near one another if they are near in t and on gaits of similar size. This relation can be formally expressed with the cartesian product of the sets $[0, T)$ and L s.t. $\ell \in L$.

[4]There exist gait families (square gaits, for example) that are only C0 continuous, and are not differentiable everywhere. These families are often not physically realistic; however, they may still be used to construct a reduced shape space provided that computations are done in regions that are locally differentiable.

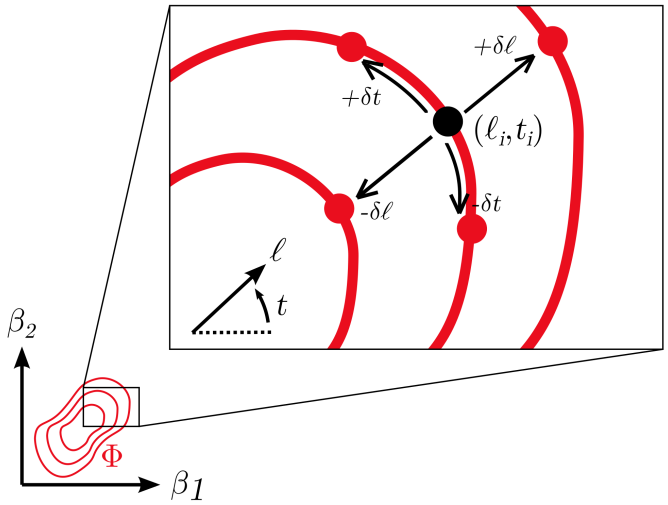

FIGURE 4.3: An arbitrary gait family Φ identifies a set of gaits $\phi \in \Phi$, ordered by gait diameter ℓ and time t. The union of Φ produces a 2D topological space, with neighborhoods defined by adjacent gaits ($\pm\delta\ell$) and time ($\pm\delta t$). The (ℓ, t) intrinsic coordinates may be reparametrized into any choice of coordinate basis; here, this is (β_1, β_2).

in B. Any chosen parametrization $(\beta_1, \beta_2) \in \mathbb{R}^2$ defines a mapping between spaces:

$$f(\beta_1, \beta_2) = (\ell, t), \tag{4.9}$$

$$\phi(\ell; t) \in S, \tag{4.10}$$

$$\phi \circ f : \mathbb{R}^2 \to B. \tag{4.11}$$

Shapes (β_1, β_2) in the reduced space S are unambiguous, so the point mapping from $\mathbb{R}^2$ to B is injective. That is, the mapping $\phi \circ f$ is an *immersion* in the shape space. In addition, the smoothness of S implies that derivatives are defined everywhere on the surface; this allows us to define the Jacobian J between S and B:

$$J : T_{(\beta_1,\beta_2)}S \to T_{\phi\circ f(\beta_1,\beta_2)}B, \tag{4.12}$$

where T_iX refers to the tangent space of X at point i.

4.3.2 Choice of Surface

Gait-joined surfaces as defined in §4.3.1 can be expressed for any smooth gait family; we are interested in a particularly effective choice of gait family to describe a reduced

shape space. Happily, recent work [27] has defined the family of step-optimal gaits for any displacement; that is, each gait in the family displaces the robot for a specific distance, with minimal cost. The ordered union of this gait family produces a reduced shape space, based on the specific objective and cost function used for gait optimization.

Of course, this is not the only candidate gait family for constructing the reduced shape space. While the surface can be constructed with other families, this choice of surface guarantees that systems can move with optimal efficiency for any displacement, as this domain contains each optimal gait in the family.

4.3.3 Implementation Details

Here, we describe the steps for implementation in code. First, we generate the surface, and find its intrinsic coordinates. Then, we generate mappings between intrinsic space and the full shape space. The two implementation components are explained here, and summarized in Algorithms 1 and 2, respectively.

Surfaces are generated by sampling optimal gaits over a finite set of displacements. Each gait is then sampled in time; the set of points across all gaits approximates the surface. Isomap [36] generates topological information (a neighborhood) for each point, and Multidimensional Scaling (MDS) is used immediately to generate intrinsic coordinates.[5] Refer to Algorithm 1 for an expression of this process.

We construct two mappings between the surface and the shape space: the point mapping defined in (4.11), and the Jacobian mapping defined in (4.12). The point mapping relies on linear interpolation between samples in extrinsic space, and is done via a Delaunay triangulation in intrinsic space. Adjacency information is pushed forward to extrinsic space to define the optimal surface. The Jacobian may be constructed immediately following the point mapping, for which we use a centered difference approximation for the derivative

[5]MDS maintains metric distances as much as possible; however, flattening of a curved surface will introduce errors in distance between some points. This may be addressed during calculations with either a pushforward of intrinsically-defined structures, or a pullback of the metric.

in intrinsic space, mapped into extrinsic space. Refer to Algorithm 2 for a summary of map construction.

Algorithm 1: Topology and intrinsic coordinates

Input: Optimal gait family Φ

Output: Intrinsic coordinates $(\beta_i, \beta_j) \in \mathbb{R}^2$

```
//get data points in all gaits
points_ND = list();
foreach φ ∈ Φ do
    points_ND.append(φ.points);
end
//generate a distance graph
D = metric_distance(points_ND);
//Isomap for intrinsic
  coordinates
points_2D = Isomap(D);
return points_ND, points_2D
```

Algorithm 2: Generation of mappings

Input: points_ND, points_2D, query_2D

Output: query_ND, query_J

```
//delaunay for interpolation
tris = delaunay(points_2D);
//point map uses triangulation
  in ND
query_ND = interpolate(tris,
  points_ND, query_2D);
//centered difference estimate
  for J
query_J =
  estimate_jacobian(point_map,
  query);
return query_ND, query_J
```

Given that only the point and Jacobian mappings are required to define the immersion, the implementation choices made in this paper may be changed at will as long as the mappings are preserved.

4.4. Results and Analysis

Here, we apply our method to the 4-link swimmer introduced in §4.2.1, both in a vacuum and and viscous fluid. We first construct an optimized surface for both environ-

ments, and visualize their immersions in the 3D shape space. Then, we show intrinsic coordinates of the optimal surface.

4.4.1 Surface Construction

We first generate a reduced shape space for the 4-link swimmer, floating in a vacuum. This system can only rotate, as it has no medium to react against; as such, we optimize for rotation only. The optimal surface, immersed in the 3D shape space, is shown in Fig. 4.4(a). As desired rotation per cycle increases, constituent gaits maintain a similar shape, and smoothly increase in diameter to incrementally capture flux of the CCF. The outer boundary is the globally most efficient gait for rotation. The location of the surface in the shape space is shown in Fig. 4.4(b), as well as the metric stretch of the space. The flux of the CCF through the surface is shown in Fig. 4.4(c). As before, the surface is oriented to maximize flux at the center, with dropoff towards the edges.

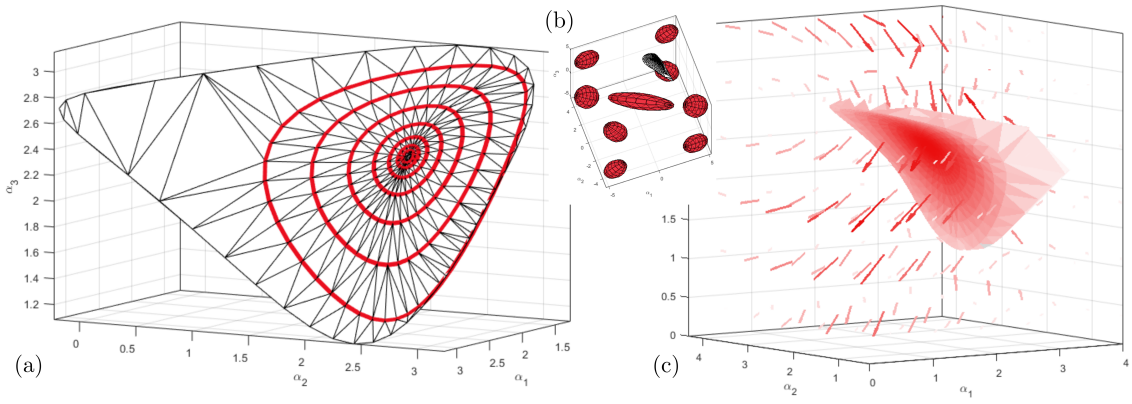

FIGURE 4.4: Rotation-optimized surface for the floating 4-link swimmer. (a) The surface is composed of optimal gaits for displacement up to the maximum-efficiency gait. (b) The Tissot indicatrix illustrates stretch of the underlying space; we have included the immersed surface, as well. (c) The CCF can be represented with a vector field in 3-space; here, we show flux of the CCF through the immersed surface.

Next, we generate a reduced shape space for the drag-dominated 4-link swimmer. This system can react against its environment and move in either the x, y, or θ directions; here, we optimize for displacement in the x direction. This optimal surface was first shown

in Fig. 4.1; the CCF flux through this surface is shown in Fig. 4.5(a). Note that flux is not absolutely maximized at the center of the surface (the CCF is not normal to the surface at the center). As shown in Fig. 4.5(b), the shape space is highly compressed at its center, normal to the surface. To capture more flux, the center gaits would necessarily rotate in this direction, and consequently be much more expensive. Instead, gaits rotate to capture more flux as their diameter increases, reflecting a decrease in relative pathlength cost towards the edges of the shape space.

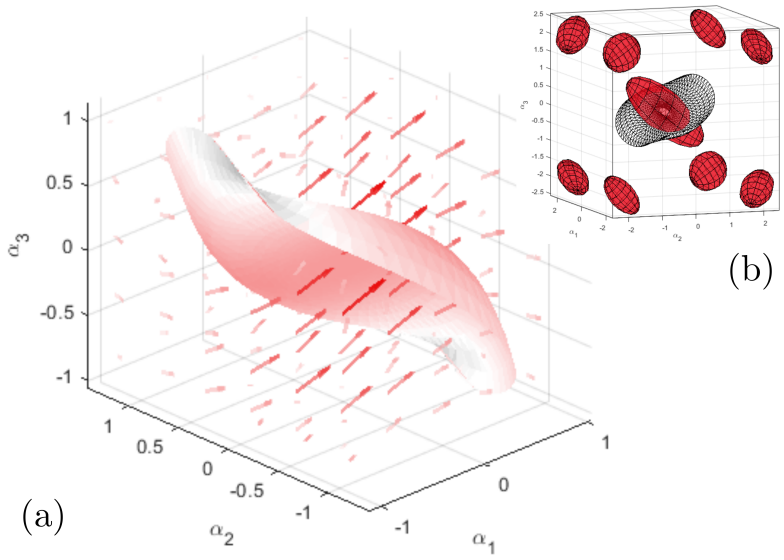

FIGURE 4.5: Optimal surface for the drag-dominated 4-link swimmer, optimized for displacement in the x direction. (a) Flux of the CCF through the surface is high, but is not necessarily normal. (b) The underlying space is highly compressed normal to the surface, making higher-flux gaits much more expensive.

4.4.2 Intrinsic Coordinates

While we can reasonably visualize the shape space of the 4-link swimmer (and any 3DoF system), the main benefit of this approach is the reduction of system behavior to 2 shape variables. Intrinsic coordinates for both swimmers were computed, and colored according to CCF flux through their corresponding surface.

The intrinsic coordinates for the drag-dominated swimmer is shown in Fig. 4.6(a);

note that CCF flux is not maximized at the center, due to the highly expensive motion required to capture additional flux. In contrast, the isolated 4-link swimmer, visualized in Fig. 4.6(b), smoothly transitions from high flux at the center to low flux at the edges. It exhibits this behavior due to the relatively low cost of gait execution in the chosen plane. Note that the high-dimensional CCF (a vector field in 3-space) has been reduced to a scalar field in the reduced space.

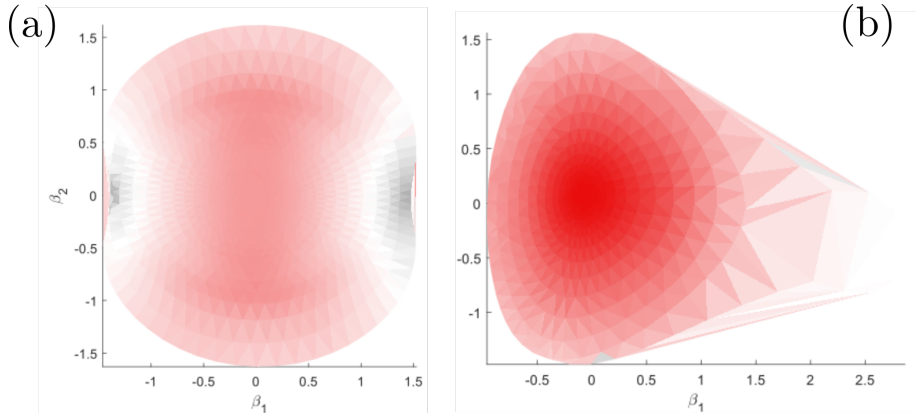

FIGURE 4.6: Intrinsic coordinates for the reduced shape space of the 4-link swimmer. (a) CCF for displacement in the x direction through a low-Re fluid. (b) CCF for rotation in a vacuum.

4.5. Conclusion

Here, we discuss important consequences of this dimensionality reduction approach for robot shape spaces, and comment on future work.

4.5.1 Consequences

Immediately, this method may be used to visualize the dynamic behavior of highly dimensional robot systems. We first identify unchanging constraints for optimization, like a single direction of interest for displacement. We then construct the reduced domain, and project important functions (in this paper, we used the CCF) onto the space. We can

then identify features of interest, and subsequently interpret them in the full shape space.

This approach my also be used to identify either unnecessary or redundant DoF. If the mappings between the reduced space and a dimension of the full shape space are relatively small, then a given dimension may not be very important for the provided objectives. In addition, if two mappings are *similar*, then the constituent DoF can likely be replaced by a single actuator.

4.5.2 Future Work

This method produces an optimal *surface* for a given choice of objective function. While further planning in the two-dimensional space could account for other goals, this reduction may be too drastic to be useful for certain applications. To address this issue, future work must explore simultaneous or iterative optimization under multiple objective functions. Instead of a 2D surface created by joining a family of gaits, the objective is to create (3+)D objects with multiple joined surfaces. These are environments that would optimize for two or more quantities, in what is effectively a n-choose-k dimensionality reduction.

Another essential next step is the application of this approach to existing robot systems, and measuring effectiveness of policies experimentally. Good candidates for such a study include aquatic robots and crawling locomotors, which behave similarly to the systems studied here.

5. GENERAL CONCLUSION AND FINAL REMARKS

The role of this work is to explore, explain, and push the boundaries of geometric mechanics. The tools used here have proven to be useful for plain analysis of mobile systems; however, their use on-hardware is limited. This is due at least in part to the complexity of this approach: the use of differential geometry in general makes system modeling more challenging. However, dynamic mobile systems are complex by nature, and require an equally rich tool to describe them. There are a number of candidate approaches, especially learning-based approaches for optimization or control. In contrast, the use of geometric mechanics provides strong understanding of the underlying dynamics and behavior of a system. In other words, we can understand *why* a particular choice of geometry, coordinate frame, or gait is the best for a particular application.

In the future, geometric mechanics may be used entirely on its own to describe and control mobile systems; it is more likely that it will be part of a stack of methods, making reliable robotics more approachable. If geometric mechanics is capable of robust modeling of kinematics and dynamics, additional methods may be used to accomplish other objectives. For example, a neural network may be given some prior knowledge of dynamics when learning a specific behavior, and could learn with more sparse information as a result. Or, we could use objective functions from geometric mechanics (like the CCF or pathlength metrics) when performing optimization, so we respect the capabilities of our robot systems.

In general, these approaches are important because they capture complexity that others do not guarantee, or miss outright. By taking the care to understand how systems behave, we can be successful where others are not. If we combine these approaches with others, we may show better results than any individual tool does on its own.

BIBLIOGRAPHY

1. S. Ramasamy and R. L. Hatton, "Soap-bubble optimization of gaits," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 1056–1062.

2. R. L. Hatton, Z. Brock, S. Chen, H. Choset, H. Faraji, R. Fu, N. Justus, and S. Ramasamy, "The geometry of optimal gaits for inertia-dominated kinematic systems," 2021.

3. R. M. Murray and S. S. Sastry, "Nonholonomic motion planning: Steering using sinusoids," *IEEE Transactions on Automatic Control*, vol. 38, no. 5, pp. 700–716, Jan 1993.

4. K. A. Morgansen, B. I. Triplett, and D. J. Klein, "Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles," *IEEE Transactions on Robotics*, vol. 23, no. 6, pp. 1184–1199, Jan 2007. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4399955

5. G. C. Walsh and S. Sastry, "On reorienting linked rigid bodies using internal motions," *Robotics and Automation, IEEE Transactions on*, vol. 11, no. 1, pp. 139–146, January 1995.

6. R. Mukherjee and D. P. Anderson, "A surface integral approach to the motion planning of nonholonomic systems," in *American Control Conference, 1993*, 1993, pp. 1816 –1823.

7. S. D. K. Kelly and R. M. Murray, "Geometric phases and robotic locomotion," *J. Robotic Systems*, vol. 12, no. 6, pp. 417–431, Jan 1995.

8. J. E. Radford and J. 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.

9. J. B. Melli, C. W. Rowley, and D. S. Rufat, "Motion planning for an articulated body in a perfect planar fluid," *SIAM Journal of Applied Dynamical Systems*, vol. 5, no. 4, pp. 650–669, November 2006.

10. E. A. Shammas, H. Choset, and A. A. Rizzi, "Geometric motion planning analysis for two classes of underactuated mechanical systems," *Int. J. of Robotics Research*, vol. 26, no. 10, pp. 1043–1073, 2007.

11. J. E. Avron and O. Raz, "A geometric theory of swimming: Purcell's swimmer and its symmetrized cousin," *New Journal of Physics*, vol. 9, no. 437, 2008.

12. R. L. Hatton and H. Choset, "Nonconservativity and noncommutativity in locomotion," *European Physical Journal Special Topics: Dynamics of Animal Systems*, vol. 224, no. 17–18, pp. 3141–3174, 2015.

13. E. M. Purcell, "Life at low Reynolds numbers," *American Journal of Physics*, vol. 45, no. 1, pp. 3–11, January 1977.

14. R. L. Hatton and H. Choset, "Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice," *International Journal of Robotics Research*, vol. 30, no. 8, pp. 988–1014, July 2011.

15. M. D. Kvalheim, B. Bittner, and S. Revzen, "Gait modeling and optimization for the perturbed stokes regime," *Nonlinear Dynamics*, vol. 97, no. 4, pp. 2249–2270, 09 2019.

16. S. Ramasamy, "Geometry of locomotion," Ph.D. dissertation, MIME, Oregon St. Univ., Corvallis OR, 2020.

17. T. Apgar, P. Clary, K. Green, A. Fern, and J. W. Hurst, "Fast online trajectory optimization for the bipedal robot cassie." in *Robotics: Science and Systems*, vol. 101, 2018, p. 14.

18. B. Katz, J. Di Carlo, and S. Kim, "Mini cheetah: A platform for pushing the limits of dynamic quadruped control," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6295–6301.

19. Boston Dynamics, "Atlas," Accessed Apr. 26, 2022]. [Online]. Available: https://www.bostondynamics.com/atlas

20. R. L. Hatton and H. Choset, "Generating gaits for snake robots by annealed chain fitting and keyframe wave extraction," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, St. Louis, MO USA, October 2009.

21. H. Marvi, C. Gong, N. Gravish, H. Astley, M. Travers, R. L. Hatton, J. R. M. III, H. Choset, D. L. Hu, and D. I. Goldman, "Sidewinding with minimal slip: snake and robot ascent of sandy slopes," *Science*, vol. 346, no. 6206, pp. 224–229, 10 October 2014.

22. R. L. Hatton and H. Choset, "Geometric swimming at low and high Reynolds numbers," *IEEE Transactions on Robotics*, vol. 29, no. 3, pp. 615–624, June 2013.

23. ——, "Kinematic cartography and the efficiency of viscous swimming," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 523–535, June 2017.

24. A. Abate, R. L. Hatton, and J. Hurst, "Passive-dynamic leg design for agile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Seattle, WA, May 2015.

25. S. Ramasamy and R. L. Hatton, "Optimal gaits for drag-dominated swimmers with passive elastic joints," *Phys. Rev. E*, vol. 103, p. 032605, Mar 2021. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevE.103.032605

26. ——, "Geometric gait optimization beyond two dimensions," in *American Control Conference (ACC), 2017*. IEEE, 2017, pp. 642–648.

27. J. Choi, C. Bass, and R. L. Hatton, "Optimal gait families using lagrange multiplier method," February 2022, submitted to IROS 2022.

28. B. Bittner, R. L. Hatton, and S. Revzen, "Geometrically optimal gaits: a data-driven approach," *Nonlinear Dynamics*, vol. 94, no. 3, pp. 1933–1948, 2018.

29. M. A. Travers, R. L. Hatton, and H. Choset, "Minimum perturbation coordinates on $so(3)$," in *Proceedings of the American Controls Conference (ACC)*, June 2013.

30. R. L. Hatton and H. Choset, "Connection vector fields for underactuated systems," in *Proceedings of the IEEE BioRobotics Conference*, October 2008, pp. 451–456.

31. C. Bass, S. Ramasamy, and R. Hatton, "Characterizing error in noncommutative geometric gait analysis," 2022. [Online]. Available: https://arxiv.org/abs/2202.10575

32. R. L. Hatton and H. Choset, "Generating gaits for snake robots: Annealed chain fitting and keyframe wave extraction," *Autonomous Robots, Special Issue on Locomotion*, vol. 28, no. 3, pp. 271–281, April 2010.

33. S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci, "Visualizing high-dimensional data: Advances in the past decade," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 3, pp. 1249–1268, 2016.

34. R. L. Hatton and H. Choset, "Approximating displacement with the body velocity integral," in *Proceedings of Robotics: Science and Systems V*, Seattle, WA USA, June 2009.

35. ——, "Connection vector fields and optimized coordinates for swimming systems at low and high Reynolds numbers," in *Proceedings of the ASME Dynamic Systems and Controls Conference (DSCC)*, Cambridge, Massachusetts, USA, Sep 2010.

36. J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.