

The Geometry of Optimal Gaits for Drag-Dominated Kinematic Systems

Suresh Ramasamy, *Student Member, IEEE*, and Ross L. Hatton , *Member, IEEE*

Abstract—In this paper, we present a set of geometric principles for understanding and optimizing the gaits of drag-dominated kinematic locomoting systems. For systems with two shape variables, the dynamics of gait optimization are analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble. The internal pressure on the gait curve is provided by the flux of the curvature of the system constraints passing through the surface bounded by the gait, and surface tension is provided by the cost associated with executing the gait, which when executed at optimal (constant-power) pacing is proportional to its pathlength measured under a Riemannian metric. We extend these principles to work on systems with three and then more than three shape variables. We demonstrate these principles on a variety of system geometries (including Purcell’s swimmer) and for optimization criteria that include maximizing displacement and efficiency of motion for both translation and turning motions. We also demonstrate how these principles can be used to simultaneously optimize a system’s gait kinematics and physical design.

Index Terms—Drag, gradient methods, low Reynolds number swimming, optimal control, optimization, robot kinematics.

I. INTRODUCTION

MANY mobile robots move by coupling cyclical internal deformations (gaits) to a continuous interaction between the robot and its environment. Because mobility is an important aspect of such robots, a key metric when evaluating design and performance of mobile robots is the efficiency of their optimal gaits. In this paper, we present a framework for understanding the geometry of optimal gaits for drag-dominated kinematic systems. This framework enables both high-level qualitative discussion of efficient gait geometries and provides a gradient for variationally solving for these efficient gaits.

In previous work, we and others in the geometric mechanics community have explored how concepts from differential geometry and geometric mechanics can be used to generate an

intuitive understanding of how systems locomote, the main factors that determine the efficiency of gaits, and what optimal gaits look like for these systems. In particular, the curvature of the system dynamics (encoded in their Lie brackets) provides useful information about classes of shape oscillations that resulted in desirable net displacements [1], [2], and we have shown that the scale of motions for which the curvature methods provide accurate information can be significantly extended by preoptimizing the choice of system coordinates [3], [4].

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

The core elements of the gait gradient are as follows.

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

As illustrated in Fig. 1, the dynamics of this optimizer are similar to the dynamics of a soap bubble, with the Lie bracket providing an internal pressure, which causes the gait cycle to expand, the metric-weighted pathlength providing a surface tension that constrains the growth of the gait, and the perimeter-balancing term providing a concentration gradient that evenly distributes waypoints along the gait.

With this basic soap-bubble framework in place for systems with two shape variables, we next extend our formulation to three dimensions by recasting some of the geometric terms, as we initially described in [6]. In particular we have the following.

- 1) We show how the constraint-curvature “area enclosure” generalizes to a flux integral through an oriented surface (more formally, an integral over a differential two-form), as illustrated in Fig. 1.

Manuscript received October 13, 2018; accepted April 23, 2019. This paper was recommended for publication by Associate Editor F. Boyer and Editor T. Murphey upon evaluation of the reviewers’ comments. This work was supported by the National Science Foundation under CMMI Grant 1462555 and Grant 1653220. (Corresponding author: Ross L. Hatton.)

The authors are with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute, Oregon State University, Corvallis, OR 97331 USA (e-mail: ramasams@oregonstate.edu; hattonr@oregonstate.edu).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TRO.2019.2915424




Fig. 1. Gaits that maximize efficiency enclose the most *curvature* of the system dynamics (measured via the curl and Lie bracket of their locomotion dynamics) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). In [5], we showed that for kinematic systems with two shape variables this process is analogous to the process by which internal pressure and surface tension combine to produce the shape and size of a soap bubble, as shown at left. More generally, the internal pressure can be thought of as being provided by the flux of the curvature passing through a surface bounded by the gait, as in the middle plot. This flux interpretation allows us to boost our geometric framework to three (right plot) or more shape variables.

- 2) We investigate how optimal gaits “bend” to exploit this orientation and maximize the flux passing through the surface.

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

As a demonstration of this approach, we use it to identify optimal gaits for a set of example systems moving in viscosity-dominated (low Reynolds number) environments. These systems include Purcell’s three-link swimmer [7] (a standard benchmark for locomotion analysis), four- and five-link swimmers, and several continuous-curvature extensions of Purcell’s swimmer with different classes and numbers of shape modes [8]. We identify optimal gaits for these systems in both the forward and turning directions, which match those found previously via raw parameter optimization in works such as [9] and [10]. The optimal gaits we find for these systems highlight the benefits of a continuous backbone. We then present how this framework can also be used for design optimization by optimizing a link length ratio and joint kinematics simultaneously.

This paper is organized as follows. In Section II, we review how the dynamics of our example systems are obtained and some of the geometric insights from our previous work. In Section III, we present our variational optimizer formulation for systems with two shape variables. In Section IV, we extend this formulation to systems with three shape variables. In Section V, we extend this formulation to systems with more than three shape variables. In Section VI, we examine how our variational formulation can be used to concurrently optimize a system’s physical structure and the gaits it executes. Section VII concludes this paper. Appendix A relates the work in this paper to prior work on these systems that made use of sub-Riemannian geometry, Appendix B provides more details on computing the system dynamics we use in our examples, and Appendix C comments on the accuracy of an approximation that enables our geometric treatment of the system dynamics.

The soap-bubble approach to gait optimization and its extensions are novel to this paper, modulo portions that appeared in conference papers [5] and [6]. We first suggested the soap-bubble framework in [5], using it to investigate optimal gaits for systems with two shape variables. We then considered the vector flux analogy for systems with three shape variables in [6]. In this paper, we consolidate, expand upon, and improve the treatments in these initial presentations. In particular, we use the theory of differential two-forms to extend our framework to systems with n shape variables and demonstrate them on several systems with four shape variables. We also expand the application of our optimizer tools, generating gaits for broader classes of systems and comparing the performance across these systems at different levels of articulation. Finally, we show how our framework can be used for simultaneous design and gait optimization.

In the context of other literature on this topic, the relationship between net displacement and area enclosed by a gait has been long recognized and commented on [1] and [11]–[16]. The results in those works, however, were either limited to small motions, or to systems that experienced pure rotation or pure translation. For systems that can have both rotational and translational velocity, the noncommutativity between these actions long prevented the use of area rules for analyzing large-amplitude gaits (even in cases where the net rotation was zero). Our work leading up to the current paper introduced a change of coordinates that mitigates the effects of noncommutativity, enabling the application of area rules to large-amplitude gaits on systems with translation/rotation freedom [3], [4], [17].

With regard to treating the gait costs as an integral over the boundary of the enclosed region, this idea appears in works such as [16] and [18]. These works were, however, limited as described in the previous paragraph; the latter work also did not include the decomposition of the perimeter cost into the length- and pacing-based contributions, which we use to gain a more completely geometric insight into the nature of optimal gaits.

Beyond combining our large-amplitude area rule with perimeter-cost formulations, in this paper, we encapsulate, abstract, and illustrate the key geometric terms involved in this




Fig. 2. Geometry and configuration variables of some of the example systems. The systems in the top row are articulated swimmers, while the bottom row consists of continuous curvature swimmers. The shape of the systems in the first columns are described by two shape variables whereas the shape of the systems in the second column are described by three shape variables.

analysis. We also explicitly extend the area rule to systems with more than two shape degrees of freedom (DoF) (which requires a novel-to-this-domain treatment of flux through two-surfaces in high-dimensional spaces) and extend our geometric approach from gait-only optimization to gait-and-morphology simultaneous optimization.

II. BACKGROUND

The geometric framework we use in this paper has its roots in works including [1] and [11]–[13], with further development in [15] and [16]. Our treatment below is condensed from a series of papers we have written for the robotics community [3], [8], [17], and at a deeper mathematical level, in [4].

For the purposes of this paper, our focus is on the geometric structure of the system dynamics. Accordingly, we work with the components of these dynamics at a relatively high level of abstraction in the equations, and present their instantiation for specific systems graphically rather than as algebraic expressions (which would run to several pages of trigonometric terms if expanded, even for the three-link swimmer). For worked examples of the construction of the dynamics of the three-link swimmer, see [3], [8], and Appendix B.

A. Geometric Locomotion Model

When analyzing a locomoting system, it is convenient to separate its configuration space Q (i.e., the space of its generalized coordinates q) into a position space G and a shape space R , such that the position $g \in G$ locates the system in the world, and the shape $r \in R$ gives the relative arrangement of the particles that compose it.¹ For example, the positions of both the articulated and continuous-curvature swimmers in Fig. 2 are the locations and orientations of their centroids and mean orientation lines, $g = (x, y, \theta) \in SE(2)$. The shape of the articulated swimmers

are parameterized by their joint angles, $r = (\alpha_1, \alpha_2)$ for the three-link swimmer and $r = (\alpha_1, \alpha_2, \alpha_3)$ for the four-link swimmer. The shape of the continuous curvature swimmers can be described by a set of modal amplitudes multiplied by the curvature modes. In the serpenoid and piecewise-continuous systems, the shape parameters α are weighting functions on curvature modes κ defined along the body, as discussed in [8].

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

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

in which $\dot{g} = g^{-1}\dot{g}$ is the body velocity of the system (i.e., $\dot{g}$ expressed in the system's local coordinates), and the *local connection* $\mathbf{A}$ acts like the Jacobian of a robotic manipulator, mapping from joint velocities to the body velocity they produce by pushing the system against its environment. Each row of $\mathbf{A}$ can be regarded as a body-coordinates local gradient of one position component with respect to the system shape. If we plot the rows of $\mathbf{A}$ as arrow fields, as in Fig. 3, this means that moving in the direction of the arrows moves the system positively in the corresponding body direction, and moving perpendicular to the arrows results in no motion in that direction [4], [17].

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

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

In line with our previous work in [23], we take the metric tensor $\mathcal{M}$ as the mapping from shape velocity to power dissipated into the surrounding medium,

$$P = \dot{r}^T \mathcal{M} \dot{r} \quad (3)$$

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

$$\tau = \mathcal{M} \dot{r} \quad (4)$$

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

As discussed in [8], [24], and [25], the length of a path under this metric describes the time it takes for a system dominated by viscous drag to follow the path at unit power (or, equivalently, the square root of the power required to execute it in unit time). Because moving with constant power is the least-costly pacing with which to execute a motion under viscous drag [24], this pathlength, thus, provides a geometric cost for the best-case execution of the kinematics in a gait cycle.

Any pacing other than constant power will make the trajectory take longer for a given average power (or increase the average

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

²This kinematic condition has been demonstrated for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [12], [19], nonholonomic constraints such as passive wheels [1], [19]–[21], and fluid interactions at the extremes of low [3], [16] and high [3], [15], [22] Reynolds numbers.




Fig. 3. Connection vector fields (a) for the three-link swimmer and (b) for the serpenoid swimmer. Scale in the vector fields has been chosen to emphasize structure, so scales in different components or systems are not comparable. Note that there is a qualitative similarity (modulo rotation) between the vector fields of the linked and serpenoid systems [3].

power required to complete the motion in a fixed time). The additional cost for a nonoptimal pacing can be represented by squaring the difference between the average and instantaneous rates at which the gait is being followed (measured as s per time), and then integrating over the time during which the gait is being executed as

$$\sigma = \int_0^{\tau_{\text{total}}} \left(\frac{s_{\text{total}}}{\tau_{\text{total}}} - \frac{d}{d\tau}(s(\tau)) \Big|_{\tau=t} \right)^2 dt \quad (5)$$

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

B. Example System Dynamics Models

In this paper, as in [3] and [8], we generate the dynamics for our example systems from a resistive force model, in which each element of the body is subject to normal and tangential drag

forces proportional to their velocities in those (local) directions.³ The normal drag coefficient is larger than the tangential component (here, by a factor of 2 : 1), corresponding to the general principle that it is harder to move a slender object in a fluid or on a surface crosswise than it is to move it along its length. We then impose a quasi-static equilibrium condition that the net drag force and moment on the system is zero at all times (treating the system as heavily overdamped, with acceleration forces much smaller than drag forces); because the drag forces are not isotropic, the system can use the angle-of-attack of its body surfaces to generate net motion.

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

$$\begin{bmatrix} F_x^b \\ F_y^b \\ F_\theta^b \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \omega(r) \begin{bmatrix} \dot{g} \\ \dot{r} \end{bmatrix} \quad (6)$$

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

By separating ω into two subblocks $\omega = [\omega_g^{3 \times 3}, \omega_r^{3 \times n}]$, it is straightforward to rearrange (6) into

$$\dot{g} = -(\omega_g^{-1} \omega_r) \dot{r} \quad (7)$$

revealing the local connection as $\mathbf{A} = \omega_g^{-1} \omega_r$.⁶ Once $\mathbf{A}$ has been found, it can be used to calculate a Riemannian metric $\mathcal{M}$ over the shape space as

$$\mathcal{M}(r) = \int_{\text{body}} J^T(r, \ell) c J(r, \ell) d\ell \quad (8)$$

where $J(r, \ell)$ is the Jacobian from shape velocity to the local velocity of each section of the body (which incorporates both $\mathbf{A}$ and the system's internal kinematics), and c is the matrix of drag coefficients.

We apply this physics model to several example geometries, which are illustrated in Fig. 2: Purcell's three-link swimmer [7], a four-link swimmer, a *serpenoid swimmer* [28], and piecewise-constant curvature swimmers with two, three, and four segments. The three-link swimmer is a useful and widely adopted [3], [9], [15], [16], [18], [29], [30] minimal example for locomotion, because its 2 DoF can be easily visualized. The serpenoid swimmer, whose shape is defined by the amplitude of sine and cosine curvature modes [23], provides an example of a 2-DoF system that has been shown to closely model how snakes and other animals use undulatory locomotion to move through

³This model is most widely associated with swimmers at low Reynolds numbers (e.g., [9]), but can also be regarded as an informative general model for systems that experience more lateral drag than longitudinal drag (e.g., [26]). Our choice of resistive force here also does not preclude the use of more detailed physical models (e.g., [27]) to construct the local connection $\mathbf{A}$.

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

⁵The expressions for the dynamics are unwieldy (running to several pages of trigonometric terms in even the simplest cases) so we do not write them out in full here. See [3] and Appendix B for a more detailed treatment of (6) and (7) in the case of the three-link swimmer, and [8] for how we build the metric tensor for that system.

⁶Details of the calculations to generate the local connection $\mathbf{A}$ and the Riemannian metric $\mathcal{M}$ for our example systems are provided in Appendix B.

the world [28]. The piecewise-constant curvature geometries are intermediaries between the discrete joints of the linked systems and the smooth traveling waves of the serpenoid systems, and correspond to system morphologies that are achievable via simple soft robotic actuators [31].

C. Gaits

Because the shape space of locomoting systems is typically bounded (e.g., because of joint limits or other restrictions on bending the body), these systems often move via gaits: cyclic changes in shape that produce characteristic net displacements. Several efforts in the geometric mechanics community [1], [2], [12], [15], [16], [19], [20] (including our own [3] and [23]) have aimed to use *curvature* of the system constraints (a measure of how “noncanceling” the system dynamics are over periodic shape changes) to understand which gaits produce useful displacements.

The core principle in these works is that because the net displacement g_ϕ over a gait cycle ϕ is the line integral of (1) along ϕ , the displacement can be approximated⁷ by a surface integral of the curvature $D(-\mathbf{A})$ of the local connection (its total Lie bracket [4]) over a surface ϕ_a bounded by the cycle,

$$g_\phi = \oint_\phi g\mathbf{A}(r) \quad (9)$$

$$\approx \iint_{\phi_a} \underbrace{-d\mathbf{A} + \sum [\mathbf{A}_i, \mathbf{A}_{j>i}]}_{D(-\mathbf{A}) \text{ (total Lie bracket)}} \quad (10)$$

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

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

$$d\mathbf{A} = \left(\frac{\partial \mathbf{A}_2}{\partial r_1} - \frac{\partial \mathbf{A}_1}{\partial r_2} \right) d\alpha_1 \wedge d\alpha_2 \quad (11)$$

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

$$[\mathbf{A}_1, \mathbf{A}_2] = \begin{bmatrix} \mathbf{A}_1^y \mathbf{A}_2^\theta - \mathbf{A}_2^y \mathbf{A}_1^\theta \\ \mathbf{A}_2^x \mathbf{A}_1^\theta - \mathbf{A}_1^x \mathbf{A}_2^\theta \\ 0 \end{bmatrix} d\alpha_1 \wedge d\alpha_2. \quad (12)$$

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

⁷This approximation (a generalized form of Stokes’ theorem) is a truncation of the Baker–Campbell–Hausdorff series for path-ordered exponentiation on a noncommutative group and closely related to the Magnus expansion [32], [33]. The accuracy of this approximation depends on the body frame chosen for the system, whose selection we discuss in [3] and [4]. In presenting this approximation, we also elide some details of exponential coordinates on Lie groups, which are also discussed in [4]. Retrospective checks on the accuracy of this approximation for some of the systems analyzed in this paper are discussed in Appendix C.




Fig. 4. Constraint curvature functions [rows of $D(-\mathbf{A})$] for (a) the three-link swimmer and (b) the serpenoid swimmer. As with the vector fields, the CCFs exhibit a qualitative similarity (modulo a rotation) between the two sets of functions.

Plotting these curvature terms as scalar functions over the shape space (as in Fig. 4) reveals the effect of gaits’ geometry on the motion they induce: Gaits that produce large net displacements in a given (x, y, θ) direction are located in strongly sign-definite regions of the corresponding $D(-\mathbf{A})$. For example, x -translation gaits encircle the center of the shape space for both the three-link and serpenoid systems, whereas y -translations or θ -rotations are produced by cycles in the corners or edges of the shape space, respectively [34].

D. Finding Optimal Gaits

Optimal gait design has a long history of research in the physics, mathematics, and engineering communities, as part of the broader field of optimal control [35], [36]. For systems of the classes, we consider here, notable contributions include those of Purcell, who introduced the three-link swimmer as a minimal template for understanding locomotion, a series of works [9], [18], [29], [30], [37] aimed at numerically optimizing the stroke pattern, the observation in [24] that the optimal pacing for the gait keeps the power dissipation constant over the cycle, the recognition that the gait optimization problem can be formulated as a variational optimization problem

with a perimeter-length cost [16], [25], and work on finding optimal gaits through two-dimensional (2-D) slices of higher dimensional shape spaces [38]. Reasonably efficient gaits are presented in [39] as natural oscillations of a locomoting system, defined as the free periodic response obtained when the damping effect is partially reduced so that the system becomes marginally stable. Optimally swimming between two points can also be considered as swimming along a sub-Riemannian geodesic [40], and a shooting-based method to identify these geodesics is presented in [41]. Numerical schemes to find these geodesics for swimming mechanisms is presented in [42] and [43].

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

- 1) The maximum-displacement gaits (the “longest strides” that the systems can take) follow the zero-contours of the constraint curvature functions, completely enclosing a sign definite region.
- 2) The most efficient gaits (where we define efficiency η as the ratio of the displacement produced by the gait to the cost of executing the gait, $\frac{g_\phi}{s}$) are contractions/straightenings of these zero-contours, as shown in Fig. 6, and correspond to the systems’ “comfortable strides”: by giving up the low-yield regions at the edges of the sign-definite regions (or crossing slightly into opposite-sign regions) the system travels a shorter distance in each cycle, but the consequently shorter perimeter length means the system can repeat the cycles more quickly at a given level of power consumption.⁸

Our definition of efficiency is equivalent to the inverse of the mechanical cost of transport used in [25] (where we prefer to work in terms of average speed achieved at a given power level instead of power required to maintain a desired average speed). Gaits that optimize our criterion also optimize Lighthill’s efficiency, which compares the power dissipated while executing the gait to the power dissipated in rigidly translating the swimmer through the fluid.

III. GRADIENT OF EFFICIENCY

To formalize the notion that efficient gaits balance the curvature they enclose against the perimeter length they require, we can write out this relationship in a differential form, providing a gradient of efficiency with respect to variations in a gait.

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

$$\nabla_p g_\phi = \mathbf{0}. \quad (13)$$

⁸This definition of efficiency is invariant with respect to time reparameterizations of the gait curve, in that it assumes that the gait curve will be reparameterized to follow at a best-case (constant-power) pacing at execution time, and in that the relative efficiency of two gaits remains the same under bulk rescaling of time. The efficiency is not, however, dimensionless with respect to system parameters such length-scale or drag coefficients, which both affect the values in the metric tensor $\mathcal{M}$ as discussed in [8].

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

$$\nabla_p \frac{g_\phi}{s} = \frac{1}{s} \nabla_p g_\phi - \frac{g_\phi}{s^2} \nabla_p s = \mathbf{0} \quad (14)$$

where the efficiency, as described in Section II-D, is taken as the displacement g_ϕ normalized by the pathlength effort required to execute the cycle [s as calculated in (2)].

For suitable seed values p_0 , solutions to (13) and (14) can, therefore, be reached by finding the respective equilibria of the dynamical systems

$$\dot{p} = \nabla_p g_\phi \quad \text{and} \quad \dot{p} = \nabla_p \frac{g_\phi}{s}. \quad (15)$$

The stable equilibria of the right-hand equation in (15) are gaits in the same “image families” as the system’s optimally efficient gaits (i.e., they follow the same curve as the optimal gait, but not necessarily at the same pacing). To construct the optimal gait, we can either optimize via (15) and then choose a constant-metric-speed parameterization, such that the pacing penalty σ from (5) goes to zero, or directly include $\nabla_p \sigma$ in our optimizer.⁹

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

$$\dot{p} = \nabla_p g_\phi - \frac{g_\phi}{s} \nabla_p s + \nabla_p \sigma \quad (16)$$

[from which we have factored out a coefficient of $1/s$ from (14)].

As illustrated in Fig. 1, this differential equation is directly analogous to the equations governing the shape of a soap bubble: $\nabla_p g_\phi$ takes the Lie bracket as an “internal pressure” seeking to expand the gait cycle to fully encircle a sign-definite region, $\nabla_p s$ is the “surface tension” that constrains the growth of the bubble, and $\nabla_p \sigma$ is the “concentration gradient” that spreads the soap over the bubble’s surface. In the following sections, we explore each of the terms in (16), discussing both their fundamental geometric definitions and how they would appear in a gradient descent based direct-transcription solver.

In our examples, we parametrize the gait as a sequence of waypoints p_i such that the gait parameters p_i explicitly define the location of the discretization points. As illustrated in Fig. 5, each waypoint p_i forms a triangle with its neighboring points and we can define a local tangent direction $e_{\parallel}$ as

$$p_{i+1} - p_{i-1} = \ell e_{\parallel} \quad (17)$$

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

We select this direct-transcription parameterization because it facilitates visualizing the workings of our optimizer (and, thus, the dynamics governing any other optimization applied to this problem). Additionally, it allows us to illustrate simultaneous optimization of the gait path and its pacing. We could also

⁹Including $\nabla_p \sigma$ in the optimizer works best for parameterizations in which $\nabla_p \sigma$ is orthogonal to $\nabla_p \frac{g_\phi}{s}$, such as waypoint-based direct transcriptions. For other parameterizations, e.g., Fourier series, the gradients may not be orthogonal and a two-step procedure of optimizing the image then the pacing will produce better results. For waypoint-based parameterizations, the $\nabla_p \sigma$ term has a secondary benefit of helping to stabilize the optimizer by maintaining an even spacing of points, and thereby preventing the formation singularities in the curve).




Fig. 5. Changes in area caused by moving in the two coordinate directions in the local frame. Moving in the tangential direction $e_{\parallel}$ produces no change in area, as the area of the triangle given by half the product of base length and height remains the same.

parametrize the gait using a Fourier series or Legendre polynomials. In this case, the pacing optimization should be done after the image of the optimal gait has been found because finding an optimal pacing can no longer be formulated as a process orthogonal to the gradient descent process for finding the image of the optimal gait.

A. Internal Pressure From the Lie Bracket

The first term in (16) $\nabla_p g_\phi$ guides the gait toward maximum-displacement cycles. By substituting the approximation from (10) into this expression as

$$\nabla_p g_\phi \approx \nabla_p \iint_{\phi_a} D(-\mathbf{A}) \quad (18)$$

and noting that variations in the gait parameters p affect the gait curve ϕ but not the system's underlying constraint $D(-\mathbf{A})$ (as $D(-\mathbf{A})$ is a property of the system and not the gait), we can convert $\nabla_p g_\phi$ into gradient of a surface integral with respect to variations in its boundary. We can then invoke a powerful geometric principle,¹⁰ which states

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

Formally, this multiplication is the *interior product*¹¹ of the boundary gradient with the integrand,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}) \quad (19)$$

which contracts $D(-\mathbf{A})$ (a *differential two-form* [45]) along $\nabla_p \phi$ to produce a differential one-form that can be integrated over ϕ . This formalism will become important in Sections IV

and V when we explore these principles on systems with more than two shape variables; for the 2-D shape spaces, we consider in this section, the interior product reduces to a simple multiplication between the outward component of $\nabla_p \phi$ and the scalar magnitude of the Lie bracket,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_{p_{\perp}} \phi) (D(-\mathbf{A})). \quad (20)$$

Implementation of the internal pressure: As illustrated in Fig. 5, the gradient of the enclosed area with respect to variations in the position of p_i , i.e., $\nabla_{p_i} \phi_a$ in the $e_{\parallel}$ and $e_{\perp}$ directions, is the change in triangle's area as p_i moves. Because the triangle's area is always one half base times height (regardless of its pitch or the ratio of its sidelengths), this gradient evaluates to

$$\nabla_{p_i} \phi_a = [e_{\parallel} \quad e_{\perp}] \begin{bmatrix} 0 \\ \ell/2 \end{bmatrix}. \quad (21)$$

Note that this term matches the right-hand side of (20), with only normal motions of the boundary affecting the enclosed area.

B. Surface Tension From the Distance Metric

The second term in (16) takes $\nabla_p s$ as a measure of how variations in the gait affect the cost of executing it, and scales this term by a factor of $\frac{g_\phi}{s}$ to compare how the return on this investment compares to the efficiency of the gait in its present state. The gradient component of this term can be related to the system's Riemannian metric by first incorporating the arclength calculation from (2) as

$$\nabla_p s = \nabla_p \oint_{\phi} \overbrace{(dr^T \mathcal{M} dr)}^{ds^2} \frac{1}{2} \quad (22)$$

¹⁰The general form of the *Leibniz integral rule* [44].

¹¹Not the inner product; the interior product contracts a two-form integrated over areas to a one-form integrated over a path by inserting a vector field as the "first" vector argument of the two-form. See [44] for more details.

and then applying standard calculus operations¹² to arrive at

$$\nabla_p s = \frac{1}{2s} \oint_{\phi} \left(2(\nabla_p dr)^T \mathcal{M} dr + dr^T (\nabla_p \mathcal{M}) dr \right) \quad (23)$$

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

The g_{ϕ}/s factor (which normalizes the scales of $\nabla_p g_{\phi}$ and $\nabla_p s$) can be calculated directly from (9) and the integral of (2). Although the calculation of g_{ϕ} could, in theory, make use of the area approximation in (10), this would be inefficient and impractical: integration of surfaces with arbitrarily complex boundaries requires significantly more computational resources than are needed for line integration around the boundary. Using the true line integral also improves the accuracy of the solution; by continuously recalibrating to the true net displacement, the algorithm avoids compounding any errors introduced by the approximation in (10).

Implementation of the surface tension: Each waypoint p_i is at the head of a vector extending from p_{i-1} , such that the dr vector and its gradient in (23) can be taken as

$$dr_i = p_i - p_{i-1} \quad (24)$$

$$\nabla_{p_i} dr_i = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad (25)$$

$$\nabla_{p_{i-1}} dr_i = \begin{bmatrix} -1 \\ -1 \end{bmatrix}. \quad (26)$$

For computational simplicity, we evaluate the metric $\mathcal{M}_i$ at the center of each segment. This point moves at the mean speed of the segment endpoints, and its gradient with respect to changes in p is the mean of its gradient at the endpoints

$$\nabla_p \mathcal{M}_i = \frac{1}{2} (\nabla_{p_{i-1}} \mathcal{M}|_{p_{i-1}} + \nabla_{p_i} \mathcal{M}|_{p_i}). \quad (27)$$

These gradient relationships are illustrated in Fig. 5, with the metric represented by its Tissot indicatrix ellipse [23].

C. Concentration Gradient From Parameterization

Once the gradient descent along (14) has given us the image of the optimal gait, a secondary gradient descent along the concentration gradient yields the optimal pacing for the gait. As discussed in Section III, if we use direct transcription to parametrize the gait, we can simultaneously optimize for the image of the optimal gait and its pacing. In this case, the third term in (16) $\nabla_p \sigma$ guides the gait toward pacings (time parameterizations) in which the cycle is executed at a constant rate of power dissipation.

Note that $\sigma = 0$ implies the mapping from τ to $r(\tau)$ gives us a constant power dissipation pacing, which we know is the optimal pacing for executing any gait from [24]. Thus, σ serves as an additional cost beyond the best pacing cost of the gait. σ compares the average and local relative rates at which the two

length coordinates are changing along the curve. Pushing the gait parameters along the gradient of this stress $\nabla_p \sigma$ brings the time parameterization into balance with the metric spacing, similarly to how the concentration gradient of soap on the surface of a bubble spreads it into a layer of uniform thickness.

The concentration gradient does not automatically appear in (13) or (14) because it is orthogonal to those optimization criteria: tangentially moving points on a curve does not change its length or the region it encloses, and so this gradient is orthogonal to the optimization gradients and acts inside their nullspace (however, see the note below on implementations of this gradient in finite gait parameterizations).

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

$$\sigma_i = ((p_{i+1} - p_i)_{\parallel} - (p_i - p_{i-1})_{\parallel})^2 \quad (28)$$

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

$$\nabla_{p_i} \sigma \propto \begin{bmatrix} e_{\parallel} & e_{\perp} \end{bmatrix} \begin{bmatrix} (p_{i+1})_{\parallel} + (p_{i-1})_{\parallel} \\ 0 \end{bmatrix} \quad (29)$$

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

Note that in the direct-transcription parameterization, each point under consideration can independently move both tangentially and perpendicularly to the gait curve and that the orthogonality between $\nabla_p \sigma$ and $\nabla_p \frac{g_{\phi}}{s}$ is, thus, preserved in the parameterization. Additionally, the spacing provided by $\nabla_p \sigma$ helps ensure the stability of the optimization by keeping the gait curve from folding over on itself and compromising the gradient calculation.

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

D. Analysis of the Purcell and Serpenoid Swimmers

We used the gradient of efficiency calculations described in the previous section to implement a gradient-descent-based optimizer in MATLAB, both directly solving the differential equation in (16) using `ode45`, and by providing (16) as the gradient for the `fmincon` optimizer using the interior-point algorithm. As expected, both implementations converged on the same solutions, with the `fmincon` implementation completing more quickly (on the order of minutes for a modern desktop computer for a gait with 100 parameters), due to its ability to take larger steps through the parameter space. Since the periodicity of the gaits lends itself to a Fourier series parametrization, for

¹²Namely: Differentiation under the integral sign, chain rule, product rule, and then exploiting the symmetry of $\mathcal{M}$ to consolidate terms.




Fig. 6. Optimal gaits found by our algorithm, overlaid onto the constraint curvature functions for their corresponding systems. Maximum-displacement gaits, which follow the zero-contours of the corresponding CCF, are indicated by dashed lines. Maximum-efficiency gaits, which are contractions of the zero-contour, are indicated by solid lines. Red regions of the CCF are positive and black regions are negative. The direction of the gait curves is matched to the sign of the region they enclose so that the displacement from each gait is positive. (a) Three-Link Swimmer- Forward Gait. (b) Serpenoid Swimmer- Forward Gait. (c) Three-Link Swimmer- Turning Gait.

the `fmincon` optimizer, we generate the direct-transcription waypoints from a lower order Fourier series parameterization. The lower order restricts the optimizer to simple gaits, preventing neighboring points from crossing each other under the discrete steps taken by `fmincon` and, thus, increasing its numerical stability. Another interesting parametrization that could be incorporated in the future is the one presented in [46], where motion of each shape variable is parametrized as the sum of a set of compactly supported bump functions added to first-order Fourier series.

We applied the optimizer to the three-link system twice, first to find the gait that maximizes the displacement in the x -direction over a single cycle, and then to find the maximum-efficiency cycle. The dashed line in Fig. 6(a) shows the gait that optimizes the maximum displacement over the cycle for the three-link swimmer. As expected, this gait follows the zero-contour of the height function. The solid line in Fig. 6(a) shows the maximum efficiency gait for this system. Because the maximum-efficiency optimizer places a cost on pathlength, this curve gives up the low-yield regions at the ends of the cycle and crosses slightly outside the zero contour. The gaits obtained via this procedure for the three-link swimmer match those obtained for maximum efficiency in [10] and, along the x direction, those in [9].

Fig. 6(b) illustrates essentially similar behavior for the serpenoid swimmer: the dashed maximum-displacement gait traces the zero contour, and the solid maximum-efficiency gait captures a more compact area within the sign-definite region. Fig. 6(c) shows the gait that produces the most cost-effective rotational motion of the three-link swimmer (the maximum-efficiency turning gait).

The low-yield region given up at the edges of the zero contour to increase efficiency depends on the anisotropy of the drag experienced by the swimmer in the lateral and longitudinal directions. Fig. 7 shows the shape spaces of serpenoid swimmers with drag ratio of (a) 2:1 and (b) 9:1 as manifolds isometrically embedded in a higher dimensional ambient space such that the cost of executing the gait is the distance traveled on the manifold while executing the gait. Such an embedding lets us accurately visualize and compare the cost of executing different

gaits. Similar visualization tools were also used in [47] to study kinematic cartography and in [48] to study 2-DoF mechanisms. The greater anisotropy would cause the embedding to have a larger curvature as seen in Fig. 7. The reduction in cost associated with the contraction of a gait is smaller on manifolds that are curved more. Thus, contractions of zero contour at a drag ratio of two reduces cost (distance traveled on the manifold) more than at a drag ratio of nine, which explains why the optimal gait is closer to the zero contour at a drag ratio of nine than at two.

IV. EXTENSION TO THREE DIMENSIONS

For systems with three shape variables, the exterior derivative and local Lie bracket terms from (10) each have three components, corresponding to the available pairs of basis vectors as

$$\begin{aligned} d\mathbf{A} = & \left(\frac{\partial \mathbf{A}_2}{\partial \alpha_1} - \frac{\partial \mathbf{A}_1}{\partial \alpha_2} \right) d\alpha_1 \wedge d\alpha_2 \\ & + \left(\frac{\partial \mathbf{A}_3}{\partial \alpha_1} - \frac{\partial \mathbf{A}_1}{\partial \alpha_3} \right) d\alpha_1 \wedge d\alpha_3 \\ & + \left(\frac{\partial \mathbf{A}_3}{\partial \alpha_2} - \frac{\partial \mathbf{A}_2}{\partial \alpha_3} \right) d\alpha_2 \wedge d\alpha_3 \end{aligned} \quad (30)$$

and

$$\sum [\mathbf{A}_i, \mathbf{A}_{j>i}] = [\mathbf{A}_1, \mathbf{A}_2] + [\mathbf{A}_1, \mathbf{A}_3] + [\mathbf{A}_2, \mathbf{A}_3]. \quad (31)$$

The surface integral in (9) is taken over an oriented surface bounded by the gait.¹³

Employing the same change of coordinates we used for the 2-D systems, we can express these two-forms with respect to a local basis in the shape space, in which $e_{\parallel}$ is tangent to the current gait, and $e_{\perp}$ is normal to the gait in its current plane.

¹³The existence of a set of such surfaces (the *Seifert surfaces* [49] of the gait) is guaranteed by knot theory; by Stokes' theorem the integral of $d\mathbf{A}$ is the same across all of these surfaces; the $\sum [\mathbf{A}_i, \mathbf{A}_{j>i}]$ integral may depend on the surface but our minimum-perturbation coordinates [17] make this term small; in any case, we are computing the change in this surface with respect to gait parameter variations, and so our algorithm does not need to explicitly identify or integrate over a specific surface.




Fig. 7. Isometric embedding of the serpenoid swimmer's shape manifold in three dimensions for a drag ratio of (a) 2:1 and (b) 9:1. The surfaces in the center and right plots are constructed such that pathlengths over the surface are equal to the metric-weighted pathlengths that capture the cost of executing a gait. The red line indicates the optimal gait when the cost is calculated using the power metric $\mathcal{M}$, and the blue line indicates the optimal gait when the power metric is replaced by an identity metric in the cost calculations (treating the cost as the simple pathlength in the parameter space). Comparing the surfaces in (a) and (b) reveals that the cost manifold for the system with a drag ratio of 9 is more sharply curved than that of the system with a drag ratio of 2. This increased curvature means that the embedding of the manifold is "steeper," and that pathlength, thus, grows more slowly with increased parameter-radius for the nine-swimmer than the two-swimmer, leading to a greater difference between the simple pathlength-optimum-stroke and the power-dissipation-optimum stroke for the nine-swimmer.

We extend this basis into three dimensions with a new vector $e_{\odot}$ that is binormal to the gait trajectory, such that displacing a gait point in this direction would rotate the local surface patch around $e_{\parallel}$.

As illustrated in Fig. 8, extending into the third dimension means that in addition to enclosing extra gait area in the $e_{\parallel} \wedge e_{\perp}$ plane by displacing gait points in the $e_{\perp}$ direction, a gait can also enclose new area in the $e_{\parallel} \wedge e_{\odot}$ plane by displacing gait points in the $e_{\odot}$ direction. The interior product from (19) now preloads the constraint curvature $D(-\mathbf{A})$ with the gradients of gait points in both the normal and binormal directions, such that the gradient of the net displacement with respect to changes in the gait parameters becomes the integral around the gait of the sum of these gradients multiplied by their respective constraint curvature components, such that

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}) \quad (32)$$

$$= \oint_{\phi} \left((\nabla_{p_{\perp}} \phi) D(-\mathbf{A})_{\parallel \perp} + (\nabla_{p_{\odot}} \phi) D(-\mathbf{A})_{\parallel \odot} \right). \quad (33)$$

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

Intuitively, this process of finding the optimal gait can be visualized as positioning the gait curve such that it maximizes the flux of a vector field corresponding to $D(-\mathbf{A})$ through a surface bounded by the curve. This vector field $\overrightarrow{D(-\mathbf{A})}$ is formed by associating each plane in the $\mathbb{R}^3$ shape space with its right-hand normal vector (noting that the e_2 basis vector is right-hand normal to the $e_3 \wedge e_1$ plane rather than the $e_1 \wedge e_3$ plane, and adjusting signs accordingly). Motions of the gait points that incorporate more sign-matched area on the $e_{\parallel} \wedge e_{\perp}$ and $e_{\parallel} \wedge e_{\odot}$ planes in the two-form interpretation serve to increase the area of the surface enclosed by the gait and to better align it with the direction of the flux field, as illustrated in Fig. 1(c).¹⁴ For systems with three shape variables, our optimizer is a generalization of [38] from optimal gaits through 2-D slices of high-dimensional spaces to arbitrary paths through these spaces.

¹⁴This vector-flux interpretation highlights the point that the *orientation* of the gait in the shape space (determined by the phase offsets between the joint oscillations) is as important as its location in the shape space (the mean and range of the joint motions) when determining the net displacement induced by the cycle.

Simply placing a gait in a region where the signs of all components of $D(-\mathbf{A})$ are the same, as suggested in [19], is neither sufficient nor necessary for generating net motion: A gait in an "all positive" region will produce no net displacement if it is orthogonal to $D(-\mathbf{A})$ (i.e., if $\overrightarrow{D(-\mathbf{A})}$ is in the plane of the gait), and a gait in a "mixed-sign" region can produce displacement proportional to the magnitude of the constraint curvature if it is aligned with $D(-\mathbf{A})$ (i.e., if $\overrightarrow{D(-\mathbf{A})}$ is normal to the surface).




Fig. 8. Changes in the area enclosed and the direction of change produced by movement along the coordinate directions in the local frame for a 3-D shape space. Moving in the tangential direction $e_{\parallel}$ produces no change in the area enclosed, as the area of the triangle given by half the product of base length and height remains the same. Moving in the $e_{\perp}$ direction increases the area along the $e_{\perp} \wedge e_{\parallel}$ plane, and moving in the e_O direction increases the area along the $e_O \wedge e_{\parallel}$ plane. Top: The forward progress of the system as it executes the optimal gait cycle.




Fig. 9. Optimal forward gaits for systems with three shape variables. (a) Four-link swimmer. The black ellipse at the center is the gait the optimizer was initialized with, and the outer red rectangle is the optimal gait found by our process. The black arrows show the vector-dual to the constraint curvature $\overrightarrow{D(-\mathbf{A})}$ at various points. This vector-dual has an approximately constant heading over the region of the shape space explored by the optimizer and the optimal gait thus approximately evolves along a plane. (b) Three-segment piecewise-constant curvature swimmer. Top: The forward progress of the systems as they execute their optimal gait cycles. Both systems move about a tenth of a body length per cycle during these gaits, with the piecewise-constant system requiring less effort for each cycle. Insets: The time-history of the three deformation modes, illustrating a 90° phase shift between the joints on the linked system, and a smaller phase shift (more tightly grouped peaks) between the segments of the piecewise-constant system.

A. Selecting a Seed Gait

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

along the zero contour of the CCF projected onto the plane most aligned with the CCF at the extremum of its absolute value.

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




Fig. 10. Projection of the CCF for the four-link swimmer onto the plane most aligned with it at its center. The blue line indicates the optimal forward gait for the four-link swimmer. (a) Projection of the optimal gait onto this plane is a contraction of the zero contour in that plane. (b) Optimal gait for the five-link swimmer only deviates slightly from this plane; hence, the zero contour of the CCF projected onto this plane provides a good seed gait for the optimizer.

plane forms the seed gait for our optimizer in systems with three shape variables.

B. Analysis of the Four-Link Swimmer

We used the optimizer to find the most efficient forward gait for a four-link (three-joint) swimmer. The optimal gait is shown in Fig. 9(a). The optimal gait for the four-link swimmer is 50% more efficient than a three-link swimmer with same total length and same drag coefficients. Three-link swimmer efficiency is 0.11, and that of the four-link swimmer is 0.16, but it is still less than the efficiency of a serpenoid swimmer with two-mode shapes (0.24).

From Fig. 10, we can see that the optimal gait for the four-link swimmer lies very close to the $\alpha_3 = -\alpha_1$ plane, such that these joints oscillate almost 180° out of phase with each other and cross each other at a zero angle. The α_2 oscillation is at approximately 90° phase difference to the other two joint motions and is slightly higher amplitude, tilting the gait within this plane.

C. Analysis of Three Segment Piecewise Continuous Curvature Swimmer

We used the optimizer to find the most efficient forward gait for a system with three piecewise-continuous curvature segments. The optimal gait is shown in Fig. 9(b). This gait is similar to the one we found for the linked system, but it has a smaller phase shift between adjacent actuators, and so is rotated around the α_3 axis by approximately 45° . The efficiency of this gait is 0.18, which is better than the four-link swimmer but still falls short of the efficiency of the serpenoid swimmer with two shape modes.

V. EXTENSION TO n DIMENSIONS

The gradient of gait displacement for systems with n -dimensional shape spaces has a similar form to the 3-D formulation in (33). The key difference is that there are now $n - 2$




Fig. 11. Number of components of the constraint curvature in n -dimensional space is equal to the number of nonparallel 2-D faces on an n -cube: one for the 2-cube, or square; three for the standard 3-cube; and six for the 4-cube (or “tesseract”). (Faces on the tesseract are considered parallel if they are formed from the same pair of basis vectors, e.g., all xw faces are parallel to each other.) For higher dimensional spaces, the number of independent 2-D faces continues to scale as $n(n - 1)/2$.

“binormal” directions¹⁵ in which gait points can be displaced, making the displacement gradient

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} \left((\nabla_{p_{\perp}} \phi) D(-\mathbf{A})_{\parallel\perp} + \sum_i^{n-2} (\nabla_{p_{\perp_i}} \phi) D(-\mathbf{A})_{\parallel\perp_i} \right). \quad (34)$$

Note that because the interior-product formulation excludes the influence any components of $D(-\mathbf{A})$ that do not include $e_{\parallel}$, the number of components in (34) goes up linearly with n , even though the number of independent planes [and, thus, components of $D(-\mathbf{A})$] scales quadratically as $\frac{n(n-1)}{2}$, as illustrated in Fig. 11.

The vector-flux analogy that we made in three dimensions becomes a “patch-flux” analogy as we move into higher dimensions: The basis areas for two-forms no longer have unique normal vectors, so we cannot directly map $D(-\mathbf{A})$ to a vector field. We can, however, still take the integral of $D(-\mathbf{A})$ over a surface with the gait as the surface boundary, and each change of the gait boundary adds an infinitesimal patch to the edge of this surface, each of which contributes to the net integral based on its alignment with the basis areas of the space and the values of $D(-\mathbf{A})$ in those basis-area directions.

A. Selecting a Seed Gait

In systems with more than three shape variables, the problem of finding a plane most aligned with the $D(-\mathbf{A})$ can be more generally restated as finding two unit vectors $v, w \in T_0 R$ ($T_0 R$ is the tangent space of R evaluated at $r = \mathbf{0}$) that maximize the value of $D(-\mathbf{A})|_{r=\mathbf{0}}(v, w)$. We shall refer to the CCF at the extremum of its absolute value $D(-\mathbf{A})|_{r=\mathbf{0}}$ as $D(-\mathbf{A})_0$. We can associate a real skew-symmetric matrix B to $D(-\mathbf{A})_0$ such

¹⁵Outside of three dimensions, “rotation around a line segment” is not a well-defined operation, so we no longer call these directions “rotational.”




Fig. 12. Optimal forward gait for the five-link swimmer. The three axes represent the first three joint angle. The value of the fourth joint angle is given by the color and thickness of the line. Thin black lines indicate sections of the gait with low values of α_4 and thick red lines indicate sections of the gait with high α_4 .

that

$$D(-\mathbf{A})_0(v, w) = v^T B w \quad \forall v, w \in T_0 R. \quad (35)$$

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

$$Bz = i\lambda z \quad (36)$$

$$B(x + iy) = -\lambda y + i\lambda x. \quad (37)$$

Therefore, $x^T B y = \lambda \|x\|_2^2$ and $y^T B x = -\lambda \|y\|_2^2$. Since B is skew symmetric, $x^T B y = -y^T B x$; therefore, $\|x\|_2 = \|y\|_2 = \frac{1}{\sqrt{2}}$ as z is a unit eigenvector.¹⁶ Therefore, if $v = \sqrt{2}x$ and $w = \sqrt{2}y$, then $D(-\mathbf{A})_0(v, w) = \lambda$. Since λ is the magnitude of the largest eigenvalue of B , this is the largest value $D(-\mathbf{A})_0$ can attain by acting on two unit vectors.

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

B. Analysis of the Five-Link Swimmer

We used the optimizer to find the most efficient forward gait for a five-link (four-joint) swimmer. The optimal gait is shown in Fig. 12. The three axes represent the first three joint angle. The value of the fourth joint angle is given by the color and

¹⁶Note that outside of the case when $D(-\mathbf{A}) = 0$ on the whole region, x cannot be equal to y as that would imply $x^T B x = -x^T B x$, hence $x = y = 0$, which would mean the largest eigenvalue of B is zero which would mean $D(-\mathbf{A})$ is zero at the extrema of its absolute value.




Fig. 13. Comparison of the efficiencies of different swimmers. Red squares represent discrete link swimmers, black circles represent piecewise continuous swimmers, and blue plus signs represent serpenoid swimmers. The number next to each symbol indicates the shape variables of the system. Displacement produced by the most efficient gait for each swimmer is given by the ordinate value, and the cost of executing the gait is given by the abscissa value of the symbol representing the swimmer. The serpenoid swimmers are the most efficient, followed by the piecewise swimmers and then the discrete link swimmers. A movie presenting a side-by-side comparison of these gaits is included in the Supplementary Material.

thickness of the line. Thin black lines indicate sections of the gait with low values of α_4 and thick red lines indicate sections of the gait with high α_4 . The optimal gait for the five-link swimmer has an efficiency of 0.21.

C. Comparison of Different Swimmers

In Fig. 13, we present a comparison of the efficiency of optimal gaits found for various systems of up to four shape dimensions. Fig. 13 shows a comparison of the cost and displacement produced by executing these gaits. For two shape modes, the serpenoid system outperforms the jointed and piecewise-constant systems, but its efficiency improves only slightly if we add a second pair of (double-spatial frequency) shape modes. The jointed and piecewise-constant systems both exhibit a pattern where moving from two to three shape DoF allows them to lower their cost of motion while increasing the net displacement, and then moving from three to four shape modes significantly increases the net displacement while increasing the cost of the motion only slightly. At any level of articulation, the piecewise-constant system outperforms the jointed system and at 4-DoF approaches the serpenoid efficiency.

VI. SIMULTANEOUS DESIGN AND GAIT OPTIMIZATION

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




Fig. 14. This figure shows the process of finding the optimal link length ratio and the maximum efficiency gait for the Purcell swimmer simultaneously. We seed our optimizer at the black gait (1), and the red gait (6) is the optimal gait. The inset figures provide a top view of the gaits at each iteration, and the numbers show the steps through which the gait and design variables evolve.

the design and gait optimization process for drag dominated systems.

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

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

By simultaneously optimizing the design and control variables, our procedure avoids having to compute the optimal gait at each intermediate set of design values, and thus requires many fewer iterations than a nested optimization scheme. If this simultaneous optimization were incorporated into a numerical optimizer without the benefit of our geometric construction of the gradient, then the design variables would compound the dimensionality of the control space, and the ensuing complexity would not be mitigated by the locality of the geometric process.

A. Identifying an Optimal Link-Length Ratio for the Purcell Swimmer

We used our optimizer to find that the optimal ratio between the length of the middle link to the length of the outer links for the Purcell swimmer is 0.80, which is close to the value of 0.75 obtained by nested optimization in [9]. The efficiency of the

optimal gait we found at the optimal link ratio is 0.12, which matches the efficiency found for the optimal gait on the 0.75-ratio system in [9] (i.e., our results are similar to previous results up to the precision of our gait-displacement ordinary differential equation solver).

The steps¹⁷ that our optimizer takes from a (deliberately nonoptimal) seed gait and equal link lengths to the optimal gait at an optimal ratio are illustrated in Fig. 14. Note that because we are simultaneously stepping along the gait and design gradients, the system does not monotonically approach the optimal link ratio, but instead moves to design variables that make the best use of the current gait geometry; as the gait geometry moves from the suboptimal ellipse to the more efficient round shape, the design variable settles to its optimal value.

VII. CONCLUSION

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

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

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

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

We demonstrated this variational principle in operation on a number of test systems in viscous-dominated environments, including Purcell’s three-link swimmer (a standard minimal template for locomotion modeling) and a serpenoid swimmer (a

¹⁷Every third iteration of `fmincon`, using our provided gradient.

model widely used in studies of animals and snake robots). In the lower dimensional cases, the optimal gaits found by our approach match those previously found by exhaustive optimizations of the gait cycles [9], and in the higher dimensional cases, the optimizer allowed us to efficiently explore a space of candidate swimming morphologies.

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

In the context of related works, our framework can be viewed as a “macroscopic” extension of the Lie-bracket control schemes in [1] and [2]. This extension is significant because it uses the geometry of the systems’ dynamics to identify the *amplitudes* of their most efficient gait cycles (in addition to the useful sets of phase-couplings identified in previous work). For our example systems swimming at low Reynolds number, the 2-D “soap-bubble” analysis can be seen as an instantiation of the boundary-value problem suggested in [25], for which we have analytically constructed a gradient from the curvature of the constraints, and the higher dimensional portions of a work to be a generalization of this principle. These extensions and generalizations of previous work in the literature have both been enabled by our calculation of a minimum-perturbation body frame [4], which significantly increases the accuracy of Lie-bracket approximations to large motions.

A key goal of our geometric study of locomotion is to make this approach feasible for robots whose dynamics are difficult to model from first principles, as in the granular systems of [26] and [50]. As they do for the ideal system models used as examples in this paper, the dimensionality extensions and explicit optimality criteria discussed here will broaden the classes of systems that we can consider in the empirical-geometric work. Additionally, the geometric structure we discuss here forms the foundation of our geometric-Floquet approach to online data-driven locomotion analysis and optimization on very high dimensional systems [46].

A line of future work that we are pursuing is to extend this optimizer to kinematic systems whose dynamics are dominated by inertial effects, including those subject to conservation of momentum and whose momentum is directed by nonholonomic constraints [51]. The cost of a gait for these types of systems depends on the acceleration through the configuration space during the gait rather than on the energy dissipated into the environment. Therefore, we expect that a cost associated with the gait’s curvature through the shape space (as well as its length) will play an important role in finding the optimal gait for these systems.

In this paper, we used our variational framework to identify optimal gaits that maximize displacement and efficiency of motion in a given direction. In future work, we plan to explore ways of handling multiobjective functions and explore ways to use the topology of constraint curvature functions to decide if gaits other than simple loops (e.g., figure-eights) would be better for some systems or objectives.

APPENDIX A

RELATION TO SUB-RIEMANNIAN GEOMETRY

In the geometric mechanics community, optimal paths between two points in the configuration spaces of swimmers in drag-dominated environments have been identified as being the sub-Riemannian geodesics joining these two points [40]. In this section, we review what sub-Riemannian geodesics are, discuss previous research aimed at finding these sub-Riemannian geodesics and present why the variational framework presented in this paper simplifies the process of finding a sub-Riemannian geodesic.

For swimmers in a drag-dominated environment, the constraints (which are completely encoded in the local connection $\mathbf{A}$) determine a vector subbundle of the tangent bundle of Q , $\mathcal{H} \subset TQ$ called the horizontal distribution. This horizontal distribution represents the set of all velocities $\dot{q}$ the swimmer can achieve. Any path that is tangent to $\mathcal{H}$ is called a horizontal path. Any absolutely continuous horizontal path that minimizes the distance (according to the power metric $\mathcal{M}$) between any points on it is called a sub-Riemannian geodesic [40].

Many methods to find these geodesics are based on the theorem proved in [52] and [53], which states that every sub-Riemannian geodesic is the cotangent projection to Q of a solution on T^*Q to the Hamiltonian differential equations for the Hamiltonian H , which is uniquely determined by the power metric $\mathcal{M}$ and the horizontal distribution $\mathcal{H}$. In general, the Hamiltonian differential equations are a set of $2n$ coupled first-order differential equations, where n is the dimension of Q . This makes calculating the sub-Riemannian geodesic a numerically expensive problem. A shooting-based method to find sub-Riemannian geodesics is presented in [41].

Gaits are closed loops in the shape space R , which can be identified as the quotient space of Q under the group action G , Q/G . With each gait $\phi : [0, 1] \rightarrow R$ and a base point $\phi(0)$, we can associate a measure of the movement along the fiber produced by executing the gait, called the representative holonomy of ϕ with respect to $\phi(0)$. The problem of finding the shortest loop with a given holonomy is called the isoholonomic problem and has been studied in detail in [40], [54], and [55]. The version of this problem most relevant to us is finding a shortest loop with a given holonomy class independent of the base point in R as studied in [56]. This shortest loop is a sub-Riemannian geodesic, and [41] presents a shooting-based method to find such a geodesic. This method was also used in [42] to identify optimal gaits for low-Re swimmers. That paper also presented analytical solutions for small shape changes that optimize efficiency. Similar schemes were used in [43] to identify optimal gaits for copepod swimmer and the Purcell swimmer. The problem addressed in this paper can be considered as an extended isoholonomic problem where amongst all the sub-Riemannian geodesics that are horizontal lifts of closed loops in the shape space and have representative holonomy in a certain desired direction, we want to find the geodesic that maximizes the efficiency η as defined in Section II-D. The numerical advantages of our method stem from the facts that the following hold.

- 1) In the minimum-perturbation coordinates, all gaits within the same image family have approximately the same holonomy. So we only have to identify the image of the gait, whose horizontal lift yields the most efficient sub-Riemannian geodesic with holonomy in a certain desired direction.
- 2) Switching to the minimum-perturbation coordinates makes the approximation in (10) valid for large angle gaits. This approximation enables quicker computation of the change in holonomy with respect to small perturbations of the gait ϕ , because the approximation makes the mapping from the image of the gait to holonomy of the gait history independent. (In the optimal coordinates, if we change one part of the gait, the contribution of the other parts toward the holonomy of the gait does not change.)

APPENDIX B

KINEMATICS OF LOW-RE SWIMMERS OBTAINED USING RESISTIVE FORCE THEORY (RFT)

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

A. Position of Links and Joints

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

$$(x, y, \theta) = \begin{bmatrix} \cos \theta & -\sin \theta & x \\ \sin \theta & \cos \theta & y \\ 0 & 0 & 1 \end{bmatrix} \quad (38)$$

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

$$h_i = \begin{bmatrix} 1 & 0 & \ell_i/2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (39)$$

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

$$a_j = \begin{bmatrix} \cos \alpha_j & -\sin \alpha_j & 0 \\ \sin \alpha_j & \cos \alpha_j & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (40)$$

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

$$g_{\lambda_{i=1}} = \text{Id} \quad (41)$$

$$g_{\lambda_{i \geq 2}} = \prod_{k=2}^i (h_{k-1})(a_k)(h_k) \quad (42)$$

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

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

$$g_{\sigma_{j=1}} = h_1 \quad (43)$$

$$g_{\sigma_{j \geq 2}} = h_1 \prod_{k=2}^j (a_{k-1})(h_k)(h_k) \quad (44)$$

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

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

$$h_{\frac{i}{j}} = (g_{\sigma_j}^{-1})(g_{\lambda_i}). \quad (45)$$

B. Velocities of Links and Joints

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

$$\dot{g}_i = \text{Ad}_{g_{\lambda_i}}^{-1} \dot{g}_1 \quad (46)$$

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

$$\text{Ad}_g^{-1} = \begin{matrix} \text{rotation} & \text{cross product} \\ \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 0 & -y \\ 0 & 1 & x \\ 0 & 0 & 1 \end{bmatrix} \end{matrix} \quad (47)$$

$$= \begin{bmatrix} \cos \theta & \sin \theta & x \sin \theta - y \cos \theta \\ -\sin \theta & \cos \theta & x \cos \theta + y \sin \theta \\ 0 & 0 & 1 \end{bmatrix}. \quad (48)$$

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

$$\dot{g}_{i>j} = \text{Ad}_{h_{\frac{i}{j}}}^{-1} \begin{bmatrix} 0 \\ 0 \\ \dot{\alpha}_j \end{bmatrix} \quad \Leftarrow \text{links distal to joint} \quad (49)$$

$$\dot{g}_{i \leq j} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \Leftarrow \text{links proximal to joint.} \quad (50)$$

If both the base link and the joints are moving, then the body velocity of the links is the sum of the contributions from the base link and joint motion. This means that we can combine the adjoint-inverse mappings from (46), (49), and (50) into a set of augmented matrices J_i that serve as Jacobians from the system's

generalized velocity (body velocity of the base link and the joint angular velocities) to the body velocities of the links,

$$\dot{g}_i = J_i \begin{bmatrix} \dot{g}_1 & \dot{\alpha}_1 & \dots & \dot{\alpha}_{i-1} & \dot{\alpha}_i & \dots & \dot{\alpha}_m \end{bmatrix}^T \quad (51)$$

where

$$J_i = \begin{bmatrix} Ad_{g_{\lambda_i}}^{-1} & Ad_{h_{\dagger}}^{-1} & \dots & Ad_{h_{\dagger-i}}^{-1} & \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \end{bmatrix} \mathbf{0}^{3 \times (n-i)} \quad (52)$$

C. Low Reynolds Number Resistive Force Model

With the kinematics of the system described in (42) and (51), we can calculate the forces on the system. At low Reynolds number, swimming bodies experience linear resistive drag, with geometry-dependent coefficients. In a full model for these forces, the drag coefficients depend on the relative proximity and orientation of the bodies, but a useful approximation to the dynamics for slender bodies is to take the drag forces on a link as being linearly related to the link's body velocity as

$$F_i^b = - \overbrace{\begin{bmatrix} L & kL \\ kL & \frac{kL^3}{12} \end{bmatrix}}^D \dot{g}_i \quad (53)$$

where k is the ratio between longitudinal and lateral drag coefficients, and the $kL^3/12$ term is the result of integrating the lateral drag on a spinning object over its length $\int_{-L/2}^{L/2} (ks) ds$. (In a more general model, the structure of (53) would be preserved, but D would become a function of α and would potentially gain off-diagonal terms.)

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

$$F_1^b = Ad_{g_{\lambda_i}}^{-*} F_i^b \quad (54)$$

which expands as

$$Ad_g^{-*} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ x \sin \theta - y \cos \theta & x \cos \theta + y \sin \theta & 1 \end{bmatrix} \quad (55)$$

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

Combining the drag matrix from (53) with the dual adjoint-inverse matrix in (55) and the Jacobians from (51), then summing over the links, produces a linear map from the system body and shape velocity to the force acting on the body, which we denote ω ,

$$F_1^b = \left(\sum_i \overbrace{Ad_{g_{\lambda_i}}^{-*} (-D) J_i}^{\omega(\alpha)} \right) \begin{bmatrix} \dot{g}_1 \\ \dot{\alpha} \end{bmatrix} \quad (56)$$

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

At low Reynolds numbers, swimmers are at quasistatic equilibrium, with the net external force equal to zero in all directions. This means that the ω term from (56) acts as a *Pfaffian constraint* on the feasible velocities for the system: for a set of $\dot{g}_1$ and $\dot{\alpha}$ velocities to be a solution to the system equations of motion, they must be in the null space of ω , i.e.,

$$[0] = \omega(\alpha) \begin{bmatrix} \dot{g}_1 \\ \dot{\alpha} \end{bmatrix}. \quad (57)$$

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

$$[0] = [\omega_g \quad \omega_\alpha] \begin{bmatrix} \dot{g}_1 \\ \dot{\alpha} \end{bmatrix}. \quad (58)$$

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

$$-\omega_g \dot{g}_1 = \omega_\alpha \dot{\alpha} \quad (59)$$

$$\dot{g}_1 = -\omega_g^{-1} \omega_\alpha \dot{\alpha} \quad (60)$$

$$\dot{g}_1 = -\mathbf{A}(\alpha) \dot{\alpha} \quad (61)$$

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

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

$$\dot{g}_i = J_i(\alpha) \begin{bmatrix} -\mathbf{A}(\alpha) \dot{\alpha} \\ \dot{\alpha} \end{bmatrix} = \left(\overbrace{J_i^\alpha(\alpha)}^{J_i(\alpha) \begin{bmatrix} -\mathbf{A}(\alpha) \\ \text{Id}^{m \times m} \end{bmatrix}} \right) \dot{\alpha} \quad (62)$$

where Id is an identity matrix.

In addition to acting as a linear map from body velocity to body force, the drag matrix D from (53) also serves as a quadratic map from body velocity of a link to the power required to maintain that velocity in the presence of viscous drag,

$$P_i = F_i^b \cdot \dot{g}_i = (D \dot{g}_i)^T \dot{g}_i = \dot{g}_i^T D \dot{g}_i. \quad (63)$$

Premultiplying and postmultiplying the drag matrix by the Jacobian in (62) pulls it back to an effective drag matrix for that link on the joint angle space,

$$P_i = \dot{\alpha}^T \left(\overbrace{J_i^{\alpha,T} D J_i^\alpha}^{D_i^\alpha} \right) \dot{\alpha}. \quad (64)$$

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

$$P = \sum_i P_i = \dot{\alpha}^T \left(\sum_i \overbrace{J_i^{\alpha,T} D J_i^\alpha}^{D^\alpha} \right) \dot{\alpha}. \quad (65)$$

As discussed in [8] and the main text, this drag matrix serves as a Riemannian metric on the joint space, such that pathlengths

$$s = \int \sqrt{d\alpha^T D^\alpha d\alpha} \quad (66)$$

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

D. Sub-Riemannian Geometry

From a sub-Riemannian perspective, the local connection and metric tensors from (61) and (65) can be calculated by first using the link Jacobians J_i to pull the drag matrices back to a metric tensor M on the full configuration space as

$$P = \begin{bmatrix} \overset{\circ}{g}_1^T & \dot{\alpha}^T \end{bmatrix} \left(\sum_i \overbrace{J_i^T D_i J_i}^M \right) \begin{bmatrix} \overset{\circ}{g}_1 \\ \dot{\alpha} \end{bmatrix}. \quad (67)$$

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

$$V \subset TQ \ni \begin{bmatrix} \overset{\circ}{g}_1 \\ \dot{\alpha} \end{bmatrix} \quad (68)$$

$$H \subset TQ \ni v^T M h = 0. \quad (69)$$

Taking the metric tensor M as having a block structure

$$M = \begin{bmatrix} M_g & M_{g\alpha} \\ M_{g\alpha}^T & M_\alpha \end{bmatrix} \quad (70)$$

(with separations corresponding to the separation between $\overset{\circ}{g}_1$ and $\dot{\alpha}$ in the generalized velocity vector), the top section of M must have same nullspace as the ω constraint calculated in (56) (and here happens to be equal to ω), because horizontal velocities must produce zero metric product with vertical velocities: the condition

$$\mathbf{0} = \begin{bmatrix} \overset{\circ}{g}_1^T & \mathbf{0} \end{bmatrix}_V M \begin{bmatrix} \overset{\circ}{g}_1 \\ \dot{\alpha} \end{bmatrix}_H \quad (71)$$

implies that

$$\mathbf{0} = \begin{bmatrix} M_g & M_{g\alpha} \end{bmatrix} \begin{bmatrix} \overset{\circ}{g}_1 \\ \dot{\alpha} \end{bmatrix}_H \quad (72)$$

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

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

space as

$$P = \dot{\alpha}^T \left(\overbrace{\begin{bmatrix} -\mathbf{A}^T & \mathbf{Id}^{m \times m} \end{bmatrix} M \begin{bmatrix} -\mathbf{A} \\ \mathbf{Id}^{m \times m} \end{bmatrix}}^{M_B} \right) \dot{\alpha} \quad (73)$$

with the result that M_B is equal to the D^α drag matrix that we calculated in (65).

E. Changing Frames

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

$$g_{\frac{\lambda_i}{0}} = (g_{\frac{0}{\lambda_1}}^{-1})(g_{\lambda_i}). \quad (74)$$

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

$$\overset{\circ}{g}_i = J_{\frac{i}{0}} \begin{bmatrix} \overset{\circ}{g}_0 \\ \dot{\alpha} \end{bmatrix}. \quad (75)$$

When calculating $J_{\frac{i}{0}}$, we need to account for the fact that changing the joint angles moves the new body frame relative to the original base link. To make the calculation, we first start by finding the Jacobian from the joint angular velocities to the new frame's body velocity, with the original body frame held fixed. This is the derivative of the g_0 frame's position with respect to the joint angles, rotated into the g_0 frame as

$$J_{\frac{0}{\alpha}} = \begin{bmatrix} \cos \theta_0 & \sin \theta_0 & 0 \\ -\sin \theta_0 & \cos \theta_0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \frac{\partial g_0(\alpha)}{\partial \alpha}. \quad (76)$$

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

$$\overset{\circ}{g}_i = \begin{bmatrix} J_{i,g} & J_{i,\alpha} \end{bmatrix} \begin{bmatrix} \overset{\circ}{g}_1 \\ \dot{\alpha} \end{bmatrix} \quad (77)$$

and use these blocks to calculate the new link velocities as

$$\overset{\circ}{g}_i = \overbrace{\begin{bmatrix} J_{\frac{i}{0}} \end{bmatrix}}^{J_{\frac{i}{0}}} \begin{bmatrix} (J_{i,g} \text{ } Ad_{g_{\frac{\lambda_1}{0}}}^{-1}) & (J_{i,\alpha} - Ad_{g_{\frac{\lambda_1}{0}}}^{-1} J_0) \end{bmatrix} \begin{bmatrix} \overset{\circ}{g}_0 \\ \dot{\alpha} \end{bmatrix} \quad (78)$$

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

rather than relative to the original base link's motion. Alternatively, the new link Jacobians can also be written directly as

$$\overset{\circ}{g}_i = \left[\overbrace{Ad_{g_{\lambda_i}^{-1}} \left[J_{i,\alpha} - Ad_{g_{\lambda_i}^{-1}} J_0 \right]}^{J_{i,0}} \right] \begin{bmatrix} \overset{\circ}{g}_0 \\ \overset{\circ}{\alpha} \end{bmatrix} \quad (79)$$

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

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

F. Continuous Backbones

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

$$h(\alpha, s) = \int_0^s \begin{bmatrix} \cos \theta(\alpha, S) & \sin \theta(\alpha, S) & 0 \\ -\sin \theta(\alpha, S) & \cos \theta(\alpha, S) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \kappa(\alpha, S) \end{bmatrix} dS. \quad (80)$$

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

$$J_\alpha(\alpha, s) = \begin{bmatrix} \cos \theta(s) & \sin \theta(s) & 0 \\ -\sin \theta(s) & \cos \theta(s) & 0 \\ 0 & 0 & 1 \end{bmatrix} \nabla_\alpha h(\alpha, s) \quad (81)$$

which can also be calculated [as a continuous analog to (52)] as

$$J_\alpha(\alpha, s) = \int_0^s Ad_{h_{\frac{s}{S}}}^{-1} \begin{bmatrix} 1 \\ 0 \\ \nabla_\alpha \kappa(\alpha, S) \end{bmatrix} dS. \quad (82)$$

Once $h(\alpha, s)$ and $J_\alpha(\alpha, s)$ have been found, $J(\alpha, s)$ can be constructed by concatenating it with the adjoint inverse from the base frame to the tangent frame at s as

$$J(\alpha, s) = \left[Ad_{h(\alpha, s)}^{-1} | J_\alpha(\alpha, s) \right] \quad (83)$$

enabling computation of the Pfaffian, shape-space drag metric, and full-configuration drag metric by substituting $\int_0^s$ for $\sum_i$ into their respective equations.

APPENDIX C

ACCURACY OF DISPLACEMENT APPROXIMATION

The geometric framework in this paper approximates the net displacement over a gait as the integral of the constraint curvature over a surface bounded by the gait. As we discussed in [3] and [4], the accuracy of this approximation degrades with gait amplitude, the rate of this degradation depends on the choice of body frame for the system, and the error in the approximation is minimized in the frame for which the norm of the connection $\mathbf{A}$ is smallest. In [3], we presented a convergence study demonstrating the accuracy of the curvature integrals for the three-link swimmer up to gait amplitudes well beyond the optimal gait amplitude.

We are actively pursuing a rigorous set of bounds on the residual error left in the constraint curvature approximation after changing into the optimized coordinates. Initial results suggest that the error introduced by this approximation is of the order θ^3 , where θ is the maximum rotation accrued while executing the gait. For the systems consider in this paper, we performed retrospective checks to ensure the error is not large. As in [3], we found no significant difference between the net displacement of the swimmers and those predicted by the surface integral, for gaits up to and significantly larger than the optimal gaits found through our approach.

REFERENCES

- [1] R. M. Murray and S. S. Sastry, "Nonholonomic motion planning: Steering using sinusoids," *IEEE Trans. Autom. Control*, vol. 38, no. 5, pp. 700–716, Jan. 1993.
- [2] K. A. Morgansen, B. I. Triplett, and D. J. Klein, "Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles," *IEEE Trans. Robot.*, vol. 23, no. 6, pp. 1184–1199, Jan. 2007. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4399955
- [3] R. L. Hatton and H. Choset, "Geometric swimming at low and high Reynolds numbers," *IEEE Trans. Robot.*, vol. 29, no. 3, pp. 615–624, Jun. 2013.
- [4] R. L. Hatton and H. Choset, "Nonconservativity and noncommutativity in locomotion," *Eur. Physical J. Special Topics, Dyn. Animal Syst.*, vol. 224, no. 17–18, pp. 3141–3174, 2015.
- [5] S. Ramasamy and R. L. Hatton, "Soap-bubble optimization of gaits," in *Proc. IEEE 55th Conf. Decis. Control*, 2016, pp. 1056–1062.
- [6] S. Ramasamy and R. L. Hatton, "Geometric gait optimization beyond two dimensions," in *Proc. Amer. Control Conf.*, 2017, pp. 642–648.
- [7] E. M. Purcell, "Life at low Reynolds numbers," *Amer. J. Phys.*, vol. 45, no. 1, pp. 3–11, Jan. 1977.
- [8] R. L. Hatton and H. Choset, "Kinematic cartography and the efficiency of viscous swimming," *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 523–535, Jun. 2017.
- [9] D. Tam and A. E. Hosoi, "Optimal stroke patterns for Purcell's three-link swimmer," *Phys. Rev. Lett.*, vol. 98, no. 6, 2007, Art. no. 068105.
- [10] R. L. Hatton, L. J. Burton, A. E. Hosoi, and H. Choset, "Geometric maneuverability, with applications to low Reynolds number swimming," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, San Francisco, CA USA, Sep. 2011, pp. 3893–3898.
- [11] A. Shapere and F. Wilczek, "Gauge kinematics of deformable bodies," *Amer. J. Phys.*, vol. 57, no. 6, pp. 514–518, Jun. 1989.
- [12] G. C. Walsh and S. Sastry, "On reorienting linked rigid bodies using internal motions," *IEEE Trans. Robot. Autom.*, vol. 11, no. 1, pp. 139–146, Jan. 1995.
- [13] J. Ostrowski and J. Burdick, "The geometric mechanics of undulatory robotic locomotion," *Int. J. Robot. Res.*, vol. 17, no. 7, pp. 683–702, 1998.
- [14] R. Mukherjee and D. P. Anderson, "A surface integral approach to the motion planning of nonholonomic systems," in *Proc. Amer. Control Conf.*, 1993, pp. 1816–1823.

- [15] J. B. Melli, C. W. Rowley, and D. S. Rufat, "Motion planning for an articulated body in a perfect planar fluid," *SIAM J. Appl. Dynamical Syst.*, vol. 5, no. 4, pp. 650–669, Nov. 2006.
- [16] J. E. Avron and O. Raz, "A geometric theory of swimming: Purcell's swimmer and its symmetrized cousin," *New J. Phys.*, vol. 9, no. 437, p. 063016, 2008.
- [17] R. L. Hatton and H. Choset, "Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice," *Int. J. Robot. Res.*, vol. 30, no. 8, pp. 988–1014, Jul. 2011.
- [18] A. DeSimone, L. Heltai, F. Alouges, and A. Lefebvre-Lepot, "Computing optimal strokes for low Reynolds number swimmers," *Natural Locomotion in Fluids and on Surfaces*. New York, NY, USA: Springer, 2012, pp. 177–184.
- [19] E. A. Shammash, H. Choset, and A. A. Rizzi, "Geometric motion planning analysis for two classes of underactuated mechanical systems," *Int. J. Robot. Res.*, vol. 26, no. 10, pp. 1043–1073, 2007.
- [20] J. P. Ostrowski and J. Burdick, "The mechanics and control of undulatory locomotion," *Int. J. Robot. Res.*, vol. 17, no. 7, pp. 683–701, Jul. 1998.
- [21] A. M. Bloch *et al.*, *Nonholonomic Mechanics and Control*. New York, NY, USA: Springer, 2003.
- [22] E. Kanso, "Swimming due to transverse shape deformations," *J. Fluid Mechanics*, vol. 631, pp. 127–148, 2009.
- [23] R. L. Hatton and H. Choset, "Kinematic cartography for locomotion at low Reynolds numbers," in *Proc. Robot., Sci. Syst. VII*, Los Angeles, CA USA, Jun. 2011.
- [24] L. Becker, S. A. Koehler, and H. A. Stone, "On self-propulsion of micro-machines at low Reynolds number: Purcell's three-link swimmer," *J. Fluid Mechanics*, vol. 490, pp. 15–35, 2003.
- [25] J. E. Avron, O. Gat, and O. Kenneth, "Optimal swimming at low Reynolds numbers," *Physical Rev. Lett.*, vol. 93, no. 18, pp. 186 001–1–186 001–r, 2004.
- [26] R. L. Hatton, H. Choset, Y. Ding, and D. I. Goldman, "Geometric visualization of self-propulsion in a complex medium," *Physical Rev. Lett.*, vol. 110, Feb. 2013, Art. no. 078101.
- [27] N. Giuliani, L. Heltai, and A. DeSimone, "Predicting and optimizing microswimmer performance from the hydrodynamics of its components: The relevance of interactions," *Soft Robot.*, vol. 5, no. 4, pp. 410–424, 2018. [Online]. Available: <https://doi.org/10.1089/soro.2017.0099>
- [28] S. Hirose, *Biologically Inspired Robots (Snake-like Locomotor and Manipulator)*. London, U.K.: Oxford Univ. Press, 1993.
- [29] P. Bettiol, B. Bonnard, L. Giralddi, P. Martinon, and J. Rouot, "The Purcell Three-link swimmer: Some geometric and numerical aspects related to periodic optimal controls," Oct. 2015. [Online]. Available: <https://hal.inria.fr/hal-01143763>
- [30] L. Giralddi, P. Martinon, and M. Zoppello, "Controllability and optimal strokes for n-link microswimmer," in *Proc. IEEE 52nd Annu. Conf. Decis. Control*, Dec. 2013, pp. 3870–3875.
- [31] C. Branyan *et al.*, "Soft snake robots: Mechanical design and geometric gait implementation," in *Proc. IEEE Int. Conf. Robot. Biomimetics*, 2017, pp. 282–289.
- [32] J. E. Radford and J. W. Burdick, "Local motion planning for nonholonomic control systems evolving on principal bundles," in *Proc. Int. Symp. Math. Theory Netw. Syst.*, Padova, Italy, 1998.
- [33] W. Magnus, "On the exponential solution of differential equations for a linear operator," *Commun. Pure Appl. Math.*, vol. VII, pp. 649–673, 1954.
- [34] L. J. Burton, R. L. Hatton, H. Choset, and A. E. Hosoi, "Two-link swimming using buoyant orientation," *Phys. Fluids*, vol. 22, Sep. 2010, Art. no. 091703.
- [35] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control, Optimization, Estimation, and Control*. New York, NY, USA: Wiley, 1975. [Online]. Available: <http://dx.doi.org/10.1002/zamm.19790590826>
- [36] J. P. Ostrowski, J. P. Desai, and V. Kumar, "Optimal gait selection for nonholonomic locomotion systems," *Int. J. Robot. Res.*, vol. 19, no. 3, pp. 225–237, 2000.
- [37] A. Shapere and F. Wilczek, "Geometry of self-propulsion at low Reynolds number," *J. Fluid Mechanics*, vol. 198, pp. 557–585, 1989.
- [38] C. Gong, D. I. Goldman, and H. Choset, "Simplifying gait design via shape basis optimization," in *Proc. Robot., Sci. Syst.*, Ann Arbor, MI, USA, Jun. 2016.
- [39] L. Zhu, Z. Chen, and T. Iwasaki, "Oscillation, orientation, and locomotion of underactuated multilink mechanical systems," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1537–1548, Sep. 2013.
- [40] R. Montgomery, *A Tour of Sub-Riemannian Geometries, Their Geodesics and Applications*. Providence, RI, USA: American Mathematical Soc., 2006, no. 91.
- [41] F. Alouges, A. DeSimone, and A. Lefebvre, "Optimal strokes for low Reynolds number swimmers: An example," *J. Nonlinear Sci.*, vol. 18, no. 3, pp. 277–302, 2008.
- [42] F. Alouges, A. DeSimone, and A. Lefebvre, "Optimal strokes for axisymmetric microswimmers," *Eur. Physical J. E*, vol. 28, no. 3, pp. 279–284, 2009.
- [43] P. Bettiol, B. Bonnard, and J. Rouot, "Optimal strokes at low Reynolds number: A geometric and numerical study of copepod and purcell swimmers," *SIAM J. Control Optim.*, vol. 56, no. 3, pp. 1794–1822, 2018.
- [44] H. Flanders, "Differentiation under the integral sign," *Amer. Math. Monthly*, vol. 80, no. 6, pp. 615–627, Jun. 1973.
- [45] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, vol. 1. Hoboken, NJ, USA: Wiley, 1963.
- [46] B. Bittner, R. L. Hatton, and S. Revzen, "Geometrically optimal gaits: A data-driven approach," *Nonlinear Dyn.*, vol. 94, pp. 1933–1948, 2018.
- [47] T. Dear, R. L. Hatton, and H. Choset, "Nonlinear dimensionality reduction for kinematic cartography with an application toward robotic locomotion," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Chicago, IL, USA, Sep. 2014, pp. 3604–3609.
- [48] G. Rodnay and E. Rimon, "Isometric visualization of configuration spaces of two-degrees-of-freedom mechanisms," *Mechanism Mach. Theory*, vol. 36, no. 4, pp. 523–545, 2001.
- [49] J. van Wijk and A. Cohen, "Visualization of Seifert surfaces," *IEEE Trans. Visualization Comput. Graph.*, vol. 12, no. 4, pp. 485–496, 2006.
- [50] J. Dai, H. Faraji, C. Gong, R. L. Hatton, D. I. Goldman, and H. Choset, "Geometric swimming on a granular surface," in *Proc. Robot., Sci. Syst. Conf.*, Ann Arbor, MI, USA, Jul. 2016, pp. 1–8.
- [51] J. Ostrowski, "Reduced equations for nonholonomic mechanical systems with dissipation," *J. Math. Phys.*, vol. 42, pp. 185–209, 1998.
- [52] C. Bär, "Carnot-caratheodory-metrien," Ph.D. dissertation, Diplomarbeit, Bonn, Univ. Bonn, Bonn, Germany, 1988.
- [53] C. Bär, "Geodesics for Carnot-caratheodory metrics," 1989, to be published.
- [54] M. J. Enos, *Dynamics and Control of Mechanical Systems: The Falling Cat and Related Problems*, vol. 1. Providence, RI, USA: American Mathematical Soc., 1993.
- [55] Z. Li and J. F. Canny, *Nonholonomic Motion Planning*, vol. 192. New York, NY, USA: Springer, 2012.
- [56] R. Montgomery, "Isoholonomic problems and some applications," *Commun. Math. Phys.*, vol. 128, no. 3, pp. 565–592, 1990.



Suresh Ramasamy (S'17) received the B.Tech. degree from the Indian Institute of Technology, Gandhinagar, India, and the M.S. degree from Carnegie Mellon University, Pittsburgh, PA, USA, both in mechanical engineering, in 2013 and 2015. He is currently working toward the Ph.D. degree in robotics and mechanical engineering at Oregon State University, Corvallis, OR, USA.

His research interests include using tools from differential geometry to understand the mechanics of locomotion, and dynamically feasible motion planning and control of nonlinear systems.



Ross L. Hatton (M'XX) received the S.B. degree from the Massachusetts Institute of Technology, Cambridge, MA, USA, and the M.S. and Ph.D. degrees from Carnegie Mellon University, Pittsburgh, PA, USA, all in mechanical engineering, in 2005, 2007, and 2011, respectively.

He is an Associate Professor of Robotics and Mechanical Engineering with Oregon State University, Corvallis, OR, USA, where he directs the Laboratory for Robotics and Applied Mechanics.

His research interests include on understanding the fundamental mechanics of locomotion and sensory perception, making advances in mathematical theory accessible to an engineering audience, and on finding abstractions that facilitate human control of unconventional locomotors.

He is a recipient of the NSF CAREER Award.