


The Geometry of Optimal Gaits for Inertia-Dominated Kinematic Systems

Ross L. Hatton , Zachary Brock, Shuoqi Chen, Howie Choset , Hossein Faraji , Ruijie Fu , Nathan Justus ,
and Suresh Ramasamy 

Abstract—Isolated mechanical systems—e.g., those floating in space, in free-fall, or on a frictionless surface—are able to achieve net rotation by cyclically changing their shape, even if they have no net angular momentum. Similarly, swimmers immersed in “perfect fluids” are able to use cyclic shape changes to both translate and rotate even if the swimmer-fluid system has no net linear or angular momentum. Finally, systems fully constrained by direct nonholonomic constraints (e.g., passive wheels) can push against these constraints to move through the world. Previous work has demonstrated that the displacement induced by these shape changes corresponds to the amount of *constraint curvature* that the gaits enclose. Properly assessing or optimizing the utility of a gait also requires considering the time or resources required to execute it: A gait that produces a small displacement per cycle, but that can be executed in a short time, may produce a faster average velocity than a gait that produces more displacement, but takes longer to complete a cycle at the same instantaneous effort. In this paper, we consider gaits under two instantaneous measures of effort. For each of these costs, we demonstrate that fixing the average instantaneous cost to a unit value allows us to transform the effort costs into time-to-execute costs for any given gait cycle. We then illustrate how the interaction between the constraint curvature and these costs leads to characteristic geometries for optimal cycles, in which the gait trajectories resemble elastic hoops distended from within by internal pressures.

Index Terms—Geometric mechanics, lie brackets, locomotion, swimming.

I. INTRODUCTION

ISOLATED mechanical systems—those floating in space, in free-fall, or on a frictionless surface—are able to achieve net rotation by cyclically changing their shape, even if they have no net angular momentum [1], [2]. Similarly, swimmers immersed

Manuscript received 20 September 2021; revised 9 January 2022; accepted 21 February 2022. Date of publication 13 May 2022; date of current version 4 October 2022. This work was supported by NSF under Grants 1462555, 1653220, 1826446, and 1935324. This paper was recommended for publication by Associate Editor P. M. Wensing and Editor P. Robuffo Giordano upon evaluation of the reviewers’ comments. (*Corresponding author: Ross L. Hatton.*)

Ross L. Hatton, Zachary Brock, Hossein Faraji, Nathan Justus, and Suresh Ramasamy are with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute, Oregon State University, Corvallis, OR 97331 USA (e-mail: ross.hatton@oregonstate.edu; zabrock@gmail.com; farajih@oregonstate.edu; justusn@oregonstate.edu; sureshr14@gmail.com).

Shuoqi Chen, Howie Choset, and Ruijie Fu are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: shuoqi@alummi.cmu.edu; choset@cs.cmu.edu; ruijief@andrew.cmu.edu).

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TRO.2022.3164595>.

Digital Object Identifier 10.1109/TRO.2022.3164595




Fig. 1. Net position displacement of an inertial-kinematic system over a gait cycle corresponds to the amount of *constraint curvature* that the gait’s trajectory encloses on the system’s shape space. The cost of executing a gait can be measured in terms of the kinetic energy or acceleration required to follow the trajectory, which correspond to the length and roundness of the trajectory over a surface shaped by the system’s inertial distribution. Gaits that maximize speed for a given energy or acceleration budget balance these two influences, with the constraint curvature acting as an “inflationary pressure” favoring large-amplitude cycles, and the pathlength and roundness terms providing a “surface tension” and “bending stiffness” that constrain the sizes and aspect ratios of the optimal gait cycles. As a general rule, the acceleration-optimal gait for a system will be both larger and rounder than the energy-optimal gait.

in “perfect fluids” are able to use cyclic shape changes to both translate and rotate even if the swimmer-fluid system has no net linear or angular momentum [3]. The operating principle in both cases is that the systems’ moments of inertia (and in the case of perfect-fluid swimming, the systems’ fluid-added masses) depend on the system shape. By moving portions of the body forward in low-inertia-configurations and backward in high-inertia configurations, systems can generate net displacements even while their total momentum remains zero, i.e., variable inertia means that conservation of momentum does not lead to conservation of position.

Given this principle, it is natural to ask “What shape trajectories best exploit the changes in inertia to generate system motion?” Answering this question involves investigating the following two subquestions:

- 1) How much displacement is achieved over a given cycle?
- 2) How much time or energy does it cost to execute a given cycle?

and then dividing the displacement by the cost to get an overall efficiency (in terms of opportunity or resources) for the motion.

In previous work, e.g., [2], [4]–[6], we and others have addressed the first question, demonstrating that the displacement of an isolated or perfect-fluid system induced by a gait cycle corresponds to the amount of *constraint curvature* in the system dynamics that the gait trajectory encloses: gaits that produce large displacements enclose strongly sign-definite regions of this curvature, and gaits that produce zero displacement enclose sign-balanced regions. As yet, however, we believe that geometric characterizations of the cost of executing different gaits under inertial dynamics (and, therefore, the optimality of such motions) have yet to be explored.

In this article, we consider two fundamental cost functions for systems with inertially dominated dynamics, corresponding, respectively, to the kinetic energy the system must have to move at a given speed and to the power losses incurred by the actuators in generating the system’s internal forces (which do not always do net work on the system). We then show that when we normalize gaits by fixing the energy or power supplied to the system to a given value, each of these cost functions can be transformed into a geometric relationship between the kinematics of the gait cycle in the shape space and the time taken to execute the gait:

- 1) When the cost is taken as the kinetic energy required for the motion, the cycle period T is proportional to the pathlength of the cycle (measured according to a metric derived from the system’s inertia), and so scales linearly with the amplitude of the gait cycle; and
- 2) When the cost is taken as the nonregenerable power consumption of the actuators, a good model for this power consumption is the square of the actuator forces, which geometrically corresponds to a squared norm of the system’s *covariant acceleration* (the acceleration relative to its unforced trajectories), weighted by the locations where the actuators are attached to the system and integrated over the gait. Under this measure of cost, the time period of a gait is proportional to the *square root* of the amplitude of the gait (because an increase in scale reduces path curvature, allowing the system to change shape more quickly at a given acceleration, and thus partially offsetting the effect that increasing inertial pathlength has on the time needed to execute the gait). Additionally, this measure of cost encourages the system to change shape more slowly in tightly curved sections of its trajectory than in straighter sections (rather than maintaining constant kinetic energy as in gaits that minimize work).

We then use these two cost functions to extend our geometric gait optimization framework (previously developed for swimmers in viscosity-dominated flows [7]–[9]) so that it provides optimal gaits for inertia-dominated systems.

Gaits optimizing speed with respect to the kinetic-energy cost function take the same “soap bubble” form as was seen for the viscous-optimized gaits (maximizing the ratio of signed constraint curvature enclosed by the gait to the pathlength of the cycle used to make this enclosure). As illustrated in Fig. 1, gaits optimizing speed with respect to the covariant acceleration cost function tend to be both larger-amplitude than

the energy-optimal motions (because the loss terms scale with the square root of amplitude rather than linearly) and rounder than the energy-optimal motions (because the need to slow down for sharp turns adds a “bending stiffness” to the gait geometry). Together, the energy-optimal and acceleration-optimal gaits bracket the optimal gaits under any cost metric that is the weighted sum of the work and loss in the actuators.

As a demonstration of the insights provided by our approach, we apply it to a range of systems, including the isolated three-link system introduced as a representative three-link model in [2], the perfect-fluid swimmer from [3], [4], two-mode continuous-curvature swimmers (which we considered in viscous flows in [9]), and generalizations of these systems to three and four shape variables (which we also considered for viscous flows in [9]).

Finally, as an extension of our core analysis, we consider optimal-gait geometry for systems—such as a three-link system with skates or passive wheels—whose locomotion is dictated by direct nonholonomic constraints, but whose dynamics within these constraints are inertially dominated.

II. BACKGROUND

The geometric framework we use in this article has its roots in works including [2], [10]–[13], with further development in [4], [14]. Our treatment below is condensed from a series of papers we have written for the robotics community [5], [9], [15], [16], and at a deeper mathematical level, in [17].

For the purposes of this article, 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 kinematics of the n -link and continuous systems, see [9, Appendix].

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 shapes 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 shapes of the continuous curvature swimmers can be described by a set of modal amplitudes multiplied by their curvature modes. In the serpenoid and piecewise-continuous systems, the shape parameters α are weighting

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*.

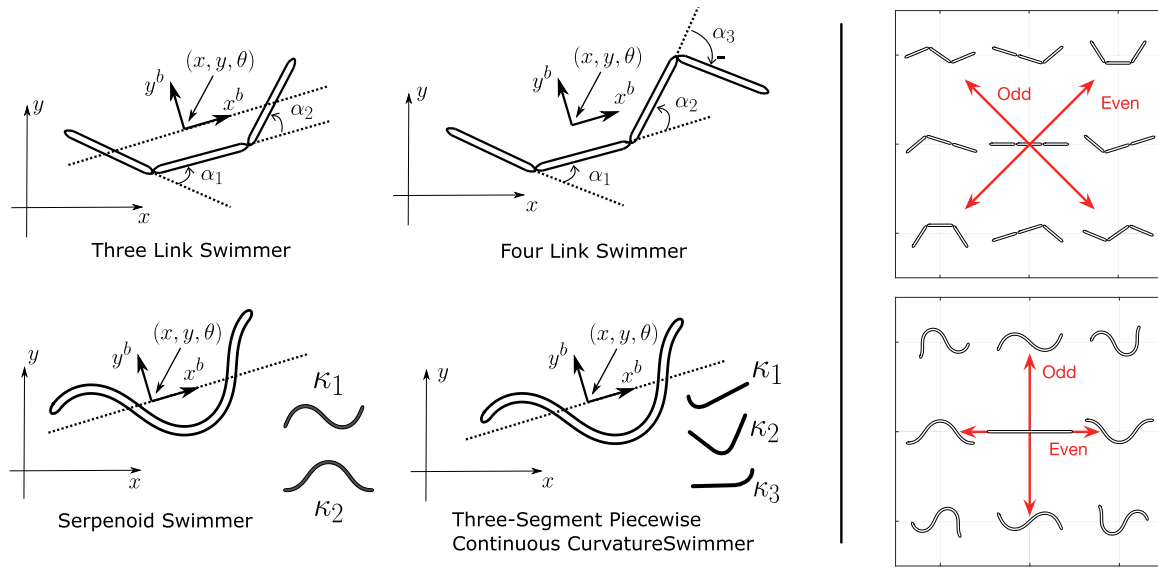


Fig. 2. Geometry and configuration variables of some of the example systems. The systems in the top row are articulated swimmers, while the bottom row consists of continuous curvature swimmers. The shape of the systems in the first column are described by two shape variables, whereas the shape of the systems in the second column are described by three shape variables. Articulated and continuous shape spaces can both be described as having “even” and “odd” axes, in which the joint angles or body curvature have even (reflection) or odd (rotational) symmetry about the system’s midpoint. These axes are illustrated in the rightmost column for systems with two shape variables; adding midpoint bending as a third shape variable creates a second “even” axis, and adding a fourth mode provides a second “odd” mode.

functions on curvature modes κ defined along the body,² as discussed in [16].³

The dynamics of an isolated-inertial or perfect-fluid locomoting system are dictated by its body-frame inertia matrix M , which relates the system’s kinetic energy KE and momentum p to its velocity as

$$\text{KE} = \frac{1}{2} \begin{bmatrix} \overset{\circ}{g}^T & \dot{r}^T \end{bmatrix} M(r) \begin{bmatrix} \overset{\circ}{g} \\ \dot{r} \end{bmatrix} \quad (1)$$

and

$$\begin{bmatrix} p_g \\ p_r \end{bmatrix} = M(r) \begin{bmatrix} \overset{\circ}{g} \\ \dot{r} \end{bmatrix} \quad (2)$$

in which $\overset{\circ}{g} = g^{-1}\dot{g}$ is the body velocity of the system (i.e., \dot{g} expressed in the system’s instantaneous body frame); r is its current shape; \dot{r} is the rate at which it is changing shape; p_g is its momentum through the world in the body-frame directions; and p_r is its momentum in the shape directions.

For an isolated n -link system, this inertia matrix can be constructed by pulling back the links’ individual inertia matrices μ_i through the Jacobians of the links,

$$M(r) = \sum_i J_i(r)^T \mu_i J_i(r) \quad (3)$$

²Body curvature is the first of four kinds of curvature that appear in this article.

³In keeping with previous convention, we use r when discussing the shape in abstract, and α for the specific parameterization of the shape space that is joint angles or modal amplitudes.

where J_i is the Jacobian that maps the body and shape velocities of the system to the body velocity of the i th link,

$$\overset{\circ}{g}_i = J_i \begin{bmatrix} \overset{\circ}{g} \\ \dot{r} \end{bmatrix}. \quad (4)$$

(Detailed calculations for this Jacobian are provided in [9].)

For a system immersed in a perfect fluid, the inertia matrix can be constructed on a similar principle, augmenting the link inertias with their fluid-added-mass matrices before pulling them back through the link Jacobians [3]. For continuously-deformable systems, the summation over links is replaced with an integral along the body,

$$M(r) = \int_{\text{body}} J^T(r, \ell) \mu(\ell) J(r, \ell) d\ell \quad (5)$$

where $\mu(\ell)$ now refers to the infinitesimal mass and moment of inertia of the portion of the body at ℓ (and can be augmented with fluid added mass as in the discrete-link case).

As described in [3] and [5], the inertia matrix M can be decomposed into blocks that, respectively, encode the system’s inertia with respect to pure position motion with fixed shape, pure shape motion at a fixed position, and the coupling terms that appear when the system is moving in both shape and position,

$$M = \begin{bmatrix} M_{gg} & M_{gr} \\ M_{rg} & M_{rr} \end{bmatrix}. \quad (6)$$

If the system starts at rest and forces are applied only through the joints (i.e., we do not “externally push” the base link), conservation of momentum means that its position-space momentum p_g remains zero for all time. Under these conditions, the upper blocks of M encode a constraint on the combinations of $\overset{\circ}{g}$ and

\dot{r} that the system can achieve,

$$p_g = \begin{bmatrix} M_{gg} & M_{gr} \end{bmatrix} \begin{bmatrix} \overset{\circ}{g} \\ \dot{r} \end{bmatrix} = \mathbf{0}. \quad (7)$$

If we take the shape velocity \dot{r} as the “known” values in this constraint equation, we can reformulate the constraint in (7) as a linear map from the shape velocity to the body velocity that puts the system at zero body-frame momentum,

$$\overset{\circ}{g} = - \overbrace{M_{gg}^{-1} M_{gr}}^{\mathbf{A}(r)} \dot{r} \quad (8)$$

in which the matrix \mathbf{A} is referred to as the *local connection* \mathbf{A} for the system, and can be thought of as the mobile-system equivalent to the Jacobian of a robotic manipulator. Each row of $-\mathbf{A}$ encodes the derivatives of one component of position (in a body frame direction) with respect to the shape components, and can be visualized as an arrow field over the shape space, as illustrated in Fig. 3 for the isolated and perfect-fluid three-link systems.

B. Gaits

Because the shape space of locomoting systems is typically bounded (e.g., by joint limits or other restrictions on bending the body), such systems often move via gaits: cyclic changes in shape that remain within the bounded region of the shape space, while producing characteristic net displacements. Gaits can be described in terms of the *path* the gait traces through the shape space, the *period* required to execute one cycle, and the *pacing* (relative timing) within the cycle.

Several efforts in the geometric mechanics community [2], [4], [11], [14], [18]–[20] (including our own [5], [9], [16]), have used the *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 (8) along ϕ , the displacement induced by a gait depends only on the gait’s path in the shape space (and not on its period or pacing). Further, the induced displacement can be approximated⁴ by a surface integral of the constraint curvature⁵ $D(-\mathbf{A})$ of the local connection (its total Lie bracket⁶) over a surface ϕ_α bounded by

⁴This 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 [21], [22]. The accuracy of this approximation depends on the body frame chosen for the system, and is most accurate for body frames at an “average” of the positions and orientations of the body segments; we discuss details of this body frame selection in [5] and [17]. In presenting this approximation, we also elide some details of exponential coordinates on Lie groups, which are also discussed in [17].

⁵Constraint curvature is the second kind of curvature to appear in this article.

⁶As discussed in [17], the system motion can be considered as a constrained motion over the full configuration space Q . The *distribution* of locally achievable motions can be identified with the vector fields $(\dot{r}_i, -g\mathbf{A}\dot{r}_i)$, and the sum of pairwise Lie brackets of these fields over Q evaluates to $(0, \mathbf{d}(-\mathbf{A}) + \sum [-\mathbf{A}_i, -\mathbf{A}_j])$, in which the latter term is a local Lie bracket on G rather than on all of Q .

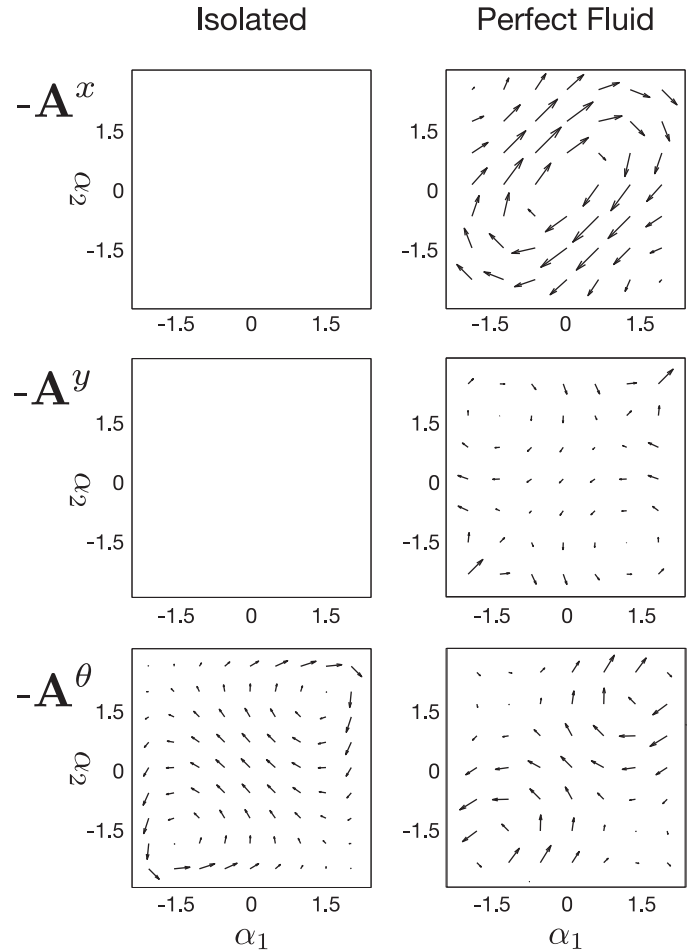


Fig. 3. Local connection \mathbf{A} for the isolated and perfect-fluid three-link systems. The x and y fields for the isolated system are zero, because the center of mass (where the body frame for this system is located) cannot move under conservation of linear momentum. Because the added mass on the links is configuration-dependent, conservation of linear momentum does not force the perfect-fluid swimmer to remain stationary, and its x and y fields are nonzero. Both systems have nonzero θ fields, because their rotational inertias are configuration dependent, so conservation of angular momentum does not prevent the system from changing orientation.

the cycle:

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

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

where $\mathbf{d}\mathbf{A}$, the exterior derivative of the local connection (its generalized rowwise 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 sequences of translations and rotations in the induced motions couple into “parallel parking” effects that contribute to the net displacement.

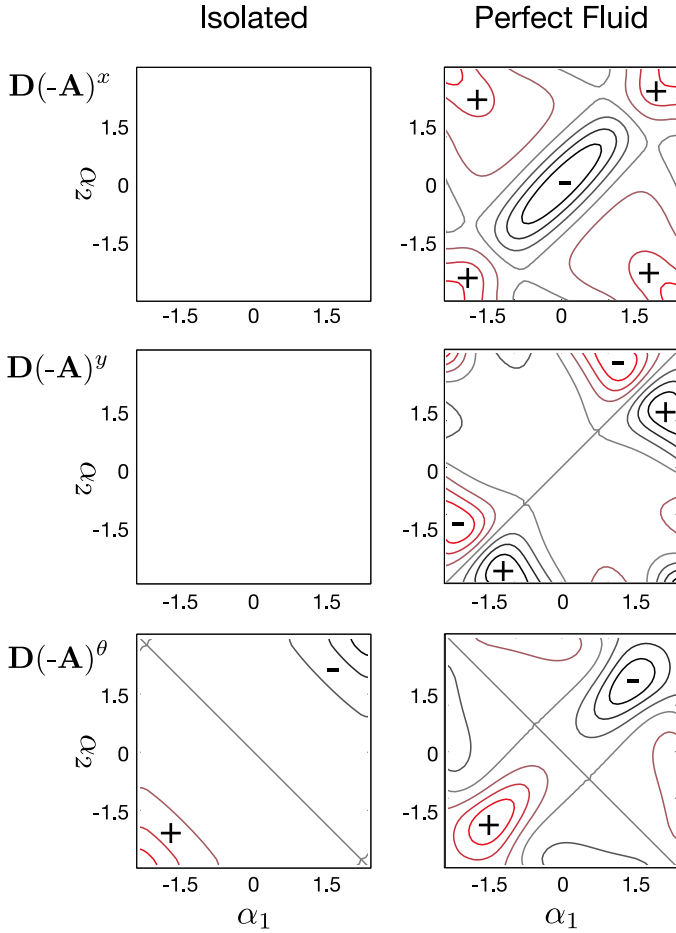


Fig. 4. Constraint curvature $\mathbf{D}(-\mathbf{A})$ for the isolated and perfect-fluid three-link systems. The θ components of the constraint curvature are the exterior derivatives (“curls”) of the $-\mathbf{A}^\theta$ fields in Fig. 3, which captures the ability of the system to make a cycle in its shape space that induces more counterclockwise body rotation than clockwise body rotation. The x and y components of the constraint curvature are the exterior derivatives of their respective $-\mathbf{A}^x$ and $-\mathbf{A}^y$ fields (capturing their ability to execute shape cycles that have more “forward” motion than “backward” motion in these body directions), augmented by the local Lie bracket term from (12), (which captures the “parallel parking” effects of mixed translation and rotation).

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

$$d\mathbf{A} = \left(\frac{\partial \mathbf{A}_2}{\partial r_1} - \frac{\partial \mathbf{A}_1}{\partial r_2} \right) dr_1 \wedge dr_2 \quad (11)$$

where \mathbf{A}_i is the i th column of the local connection (corresponding to the i th shape variable), and the local Lie bracket term for planar translation and rotation evaluates as

$$[\mathbf{A}_1, \mathbf{A}_2] = \begin{bmatrix} \mathbf{A}_1^y \mathbf{A}_2^\theta - \mathbf{A}_2^y \mathbf{A}_1^\theta \\ \mathbf{A}_2^x \mathbf{A}_1^\theta - \mathbf{A}_1^x \mathbf{A}_2^\theta \\ 0 \end{bmatrix} dr_1 \wedge dr_2. \quad (12)$$

In both cases, the wedge product $dr_1 \wedge dr_2$ indicates the oriented differential area basis in the shape space.

Plotting these curvature terms as scalar functions over the shape space (as in Fig. 4) reveals the effect of gaits’ geometry on the motions 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 $\mathbf{D}(-\mathbf{A})$ constraint curvature functions (CCFs). For example, θ rotations of both the isolated and perfect-fluid three-link systems are produced by cycles in the corners of the shape space, whereas cycles centered in the shape space produce net x translations of the perfect-fluid swimmers.

III. INERTIAL DYNAMICS

The linear map from shape velocity to body velocity in (8) means that the system’s kinetic energy, expressed as a function of the body and shape velocity in (1), can be expressed entirely as a function of the shape velocity,

$$\begin{aligned} \text{KE} &= \frac{1}{2} \dot{r}^T \begin{bmatrix} -\mathbf{A}(r)^T & \text{Id} \end{bmatrix} \underbrace{\begin{bmatrix} M_{gg} & M_{gr} \\ M_{rg} & M_{rr} \end{bmatrix}}_{M(r)} \begin{bmatrix} -\mathbf{A}(r) \\ \text{Id} \end{bmatrix} \dot{r} \quad (13) \\ &= \frac{1}{2} \dot{r}^T M_r(r) \dot{r}. \quad (14) \end{aligned}$$

Note that M_r , the *reduced inertia matrix* for the system, is distinct from the M_{rr} block of M (which encodes the shape-space inertia for the case where the body frame is fixed, whereas M_r takes the body frame as moving according to the relationship encoded in $-\mathbf{A}$).

Because M_r encodes the system’s full inertial information as a structure on the shape space, it allows us to evaluate the dynamics of the system entirely on the shape space, in what is called a *sub-Riemannian* approach [9], [23]. We briefly review here some key ideas in evaluating the motion of systems with second-order dynamics, with a focus on the geometry underlying the dynamics of these systems.

A. Shape-Space Dynamics

If we take the kinetic energy expression in (14) as defining a Lagrangian for the system, the Euler–Lagrange equations dictate a relationship between shape-space forces τ and motion through the shape space

$$\tau = M_r(r) \ddot{r} + C_r(r, \dot{r}) \quad (15)$$

in which the vector C_r (which encodes the centrifugal and Coriolis forces acting on the system) is calculated from the derivatives of the mass matrix with respect to the shape variables as

$$C_r(r, \dot{r}) = \left(\sum_{i=1}^d \frac{\partial M_r(r)}{\partial r_i} \dot{r}^i \right) \dot{r} - \frac{1}{2} \begin{bmatrix} \dot{r}^T \frac{\partial M_r(r)}{\partial r_1} \dot{r} \\ \vdots \\ \dot{r}^T \frac{\partial M_r(r)}{\partial r_d} \dot{r} \end{bmatrix}. \quad (16)$$

In some cases (including later sections of this article), it is useful to factor the right-hand side of (15) as

$$\tau = M_r(r) \overbrace{(\ddot{r} + M_r^{-1}(r) C_r(r, \dot{r}))}^{a_{\text{cov}}} \quad (17)$$

in which a_{cov} , the system’s *covariant acceleration*, describes the rate at which it is accelerating relative to its natural (unforced)

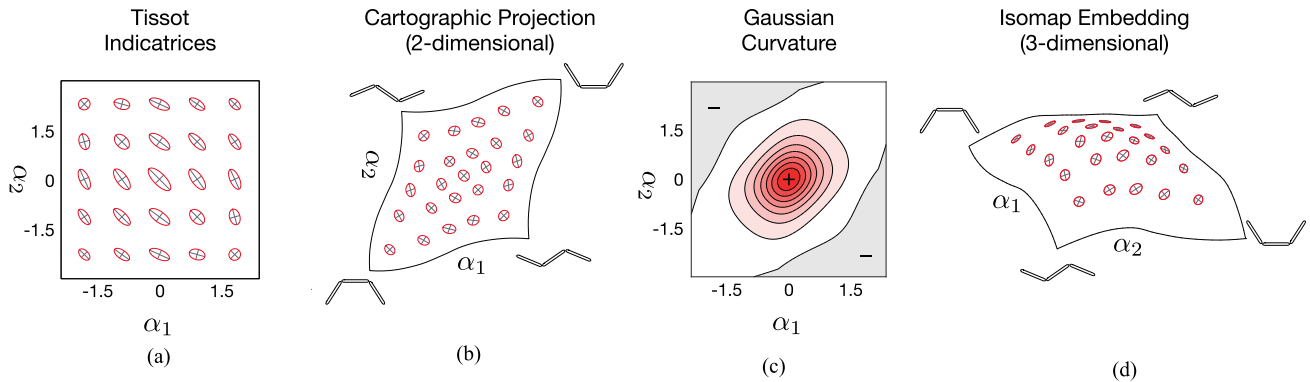


Fig. 5. Inertia matrices for our systems act as *Riemannian metrics* on their shape spaces. (a) In coordinates, these metrics can be visualized via their Tissot indicatrices (the sets of velocity vectors at each point that produce unit kinetic energy under the metric). (b) To (approximately) recover the geometry of the corresponding inertial manifold, we can use algorithms such as cartographic projection [16] to stretch the shape space so that the indicatrices become as close to uniform circles as is possible, while maintaining the continuity of the embedding. The inertial metric illustrated here is specifically that of the three-link isolated system, but it is qualitatively the same as those of the other systems we consider here: the true inertial geometry is stretched relative to the system coordinates along the even axis of the shape space. (c) The structure of the metric means that the inertial manifold has positive Gaussian curvature (i.e., it is domed/cupped rather than saddle-shaped) at the center, and becomes flat (zero curvature) to slightly saddle-shaped (negative curvature) toward the edges. (d) This domed structure can be partially visualized by using an algorithm such as Isomap [24] to approximate an isometric embedding of the manifold in 3-D space.

trajectories.⁷ The magnitude of the covariant acceleration, taken with respect to the inertia matrix M_r as

$$\|a_{\text{cov}}\| = \sqrt{a_{\text{cov}}^T M_r a_{\text{cov}}} \quad (18)$$

is equal to the total (mass-weighted) acceleration of the particles making up the system *that is not due to the constraint forces*, and, thus, directly proportional to the total force that must be actively applied to the particles to follow the trajectory.

Taking the magnitude of a_{cov} with respect to the *square* of the inertia matrix (i.e., taking the “ M^2 -norm of the acceleration”), returns a value equal to the Euclidean norm of the shape forces,

$$\|a_{\text{cov}}\|_{\tau} = \sqrt{a_{\text{cov}}^T M_r^2 a_{\text{cov}}} = \sqrt{\tau^T \tau}. \quad (19)$$

If the shape space parameterization is chosen such that the shape parameters correspond to the directly actuated degrees of freedom (e.g., using a joint-angle parameterization for an articulated chain whose joints are individually driven by motors), then $\tau^T \tau$ is the sum of squared actuator forces, and $\|a_{\text{cov}}\|_{\tau}^2$ then corresponds to the effort the system actually has to exert to achieve the motion.^{8, 9}

⁷The $M_r^{-1} C_r$ term evaluates to the *Christoffel symbols* associated with M_r . We prefer this expression over the formula that directly encodes the Christoffel symbols because it more directly illustrates how the system inertia contributes to the final expression of the dynamics.

⁸The magnitude of forces applied to the particles is different from the magnitude of the actuator forces because the latter value accounts for the leverage that the actuators have on the masses. Taking the τ -norm of a_{cov} is *not* equivalent to replacing the metric M_r with a new metric: the covariant acceleration is still calculated with respect to M_r , as in (17), and the M_r^2 term in (19) specifically corresponds to the presence of a *cometric* that is not dual to the metric, and which calls out the actuator forces as being a better measure of effort than the net forces acting on the particles.

⁹Note that although the standard measure of covariant acceleration in (18) is invariant under changes of coordinates, the τ -norm in (19) depends on the choice of shape coordinates (and is physically meaningful because the choice of coordinates is tied here to the physical placement of the actuators on the system, and, thus, the leverage that they have on the system masses).

As we discuss in Section IV, these shape-space dynamics can be combined with the constraint-curvature analysis in Section II-B, which predicts the net displacement induced by a gait from a structure that is also defined on the shape space, to perform a full cost-and-displacement analysis using only shape space structures.

B. Inertial Geometry

The reduced inertia matrix acts as a *Riemannian metric* on the shape space, defining a weighted two-norm for velocity vectors on the shape space such that the inertially weighted speed of the system is the square root of twice the system’s kinetic energy,

$$\|\dot{r}\|_{M_r} = \sqrt{\dot{r}^T M_r(r) \dot{r}} = \sqrt{2\text{KE}}. \quad (20)$$

This norm allows us to visualize the structure of the reduced inertia matrix by constructing a set of *Tissot indicatrices* [16], [25], as illustrated in Fig. 5. The indicatrix at each point in the shape space is the set of \dot{r} vectors with unit norm ($\|\dot{r}\|_{M_r} = 1$) at that point, which form an ellipse in the corresponding tangent space. The short axes of the indicatrices correspond to configurations and directions in which the system is “heavy” with respect to the joints, such that less joint motion is required to achieve a given kinetic energy. Conversely, the long axes of the indicatrices correspond to configurations and directions in which the system is “light” with respect to the joints, such that more joint motion is required to produce a given kinetic energy.

For example, the shape-position velocity coupling means that at a given $\dot{\alpha}$ joint velocity, it takes less kinetic energy for a three-link system to move along the “odd” (S-shaped) axis of the shape space than it takes to move along the “even” (C-shaped) axis. As illustrated in Fig. 5(a), the Tissot indicatrices (sets of unit-norm velocities) are longer $\dot{\alpha}$ vectors in directions along the odd axis than they are along the even axis.

Some of the true (inertial) geometry of the system can be captured by “stretching” the shape space [16] so that the Tissot

indicatrices become closer to being circles (much in the same way that a good map projection can reduce distortion of map features). For the three-link system, this means stretching along the even axis of the shape space, as illustrated in Fig. 5(b), revealing that pairs of C-shapes with a given magnitude of α are inertially further apart from each other than are pairs of S-shapes with corresponding joint angles.

In general, the true inertial geometry is *curved*, and cannot be completely recovered simply by stretching the space (much as the curvature of the Earth means that no flat map can ever be completely distortion-free). The *Gaussian curvature*¹⁰ of the inertial geometry can be calculated from M_r and its derivatives across the shape space via the Brioschi formula [26]. The Gaussian curvature of the three-link system, illustrated in Fig. 5(b) is mostly positive (meaning that the system's inertial geometry is cupped or domed), except for some negative regions at the edges of the space under consideration (meaning that the true geometry is saddle-shaped or wrinkled in these regions). Some of this curvature can be directly visualized by using an algorithm, such as Isomap [24] to approximate an isometric embedding of the inertial manifold into 3-D space, as in Fig. 5(c); in most cases, however, more than one extra dimension is required to exactly represent the geometry.

In the subsequent sections of this article, all of our calculations will be in the natural shape coordinates for each system, using M_r as a set of weights on shape velocities and accelerations. The manifold visualizations, however, play a key role in understanding the results of these calculations and the features we expect to see in the results:

The dynamics of the inertia-weighted system on the shape space are equivalent to those of a unit point mass constrained to move over the inertia-defined surface.

Recognizing this equivalence provides a context and vocabulary for discussing the geometry of the systems' optimal trajectories.

C. Geometry of the Shape-Space Dynamics

An intuitive understanding of the geometry of the shape-space dynamics can be achieved by transforming the system's velocity and covariant acceleration into a local basis constructed such that it is orthonormal with respect to the reduced mass matrix M_r and that the first vector in the basis is in the same direction as \dot{r} . In this "primed" frame, the system velocity takes the form

$$\dot{r}' = \begin{bmatrix} v \\ \mathbf{0} \end{bmatrix} \quad (21)$$

where $v = \|\dot{r}'\|_{M_r} = \sqrt{2\text{KE}}$ is both the inertial speed of the system and the velocity of the corresponding unit point mass traveling over the inertia-defined surface with the same \dot{r} parameter velocity.

¹⁰Gaussian curvature of the inertial geometry is the third kind of curvature to appear in this article.

The inertia-weighted pathlength of a trajectory through the shape space (calculated with respect to the system coordinates) can then be thought of as being equivalent to both the pathlength of the corresponding trajectory embedded into the inertia-defined surface (the integral of v),

$$S = \int_{\phi} \sqrt{\dot{r}^T M_r(r) \dot{r}} dt = \int_{\phi} v(t) dt = \int_{\phi} \sqrt{2\text{KE}(t)} dt \quad (22)$$

where the inertia-weighted expression lets us avoid explicitly constructing the basis for the primed frame. This inertial pathlength integral can be written without reference to time as

$$S = \int_{\phi} \sqrt{dr^T M_r(r) dr} \quad (23)$$

meaning that the integrated square root of kinetic energy is a property of the trajectory's path, and is specifically independent of both the period and pacing with which it is followed.

The shortest inertial paths between points in the shape space are *geodesic paths* for the system. The geodesic paths are also the generalizations of "straight paths" within the inertial surface, and the system's natural (unforced trajectories), or *geodesics*, follow the geodesic paths at constant inertial speed.

If the the metric-orthonormal, path-aligned basis is additionally constructed such that the covariant acceleration lies in the plane formed by the first and second basis vectors, the covariant acceleration takes the form

$$a'_{\text{cov}} = \begin{bmatrix} \dot{v} \\ \kappa v^2 \\ \mathbf{0} \end{bmatrix} \quad (24)$$

where \dot{v} is the rate at which the inertial speed is changing and κ is the curvature of the trajectory *within the surface*.¹¹ Together, these components describe the extent to which the system is deviating from an unforced trajectory.

The squared norm of the covariant acceleration, taken in bases orthonormal with respect to the metric tensor, is equal the squared sum of the speed- and direction-change components of the acceleration in the orthonormal bases,

$$\|a_{\text{cov}}\|_{M_r}^2 = a_{\text{cov}}^T M_r a_{\text{cov}} = a'_{\text{cov}} \cdot a'_{\text{cov}} \quad (25)$$

$$= \dot{v}^2 + (\kappa v^2)^2. \quad (26)$$

This norm of the acceleration corresponds to the amount of force that must be applied *tangent to the surface* to push the unit point mass along its trajectory, or equivalently to the magnitude of the vector-sum of nonconstraint forces that must be applied to the particles in the system.

In metric-orthonormal coordinates, the squared norm of forces appears as a weighted inner product

$$\|a_{\text{cov}}\|_{\tau}^2 = \tau^T \tau = a_{\text{cov}}^T \overbrace{\begin{bmatrix} b_{\dot{v}} & b_{v\dot{v}} \\ b_{v\dot{v}} & b_v \end{bmatrix}}^{M_{\tau}} a'_{\text{cov}} \quad (27)$$

¹¹Trajectory curvature is the fourth and final kind of curvature to appear in this article.

$$= b_v \dot{v}^2 + 2b_{v\dot{v}}(\dot{v}\kappa v^2) + b_v(\kappa v^2)^2 \quad (28)$$

in which M_T can be constructed as $J_\perp J_\perp^T$, where J_\perp is the Jacobian mapping from coordinate bases to the path-aligned metric-orthonormal bases. These weighting terms capture the property that the actuator forces are not expressed in bases orthonormal with respect to M (equivalently, are not applied directly at the particles), but instead are exerted at specific points on the mechanism, with varying amounts of “leverage” on the system particles.

IV. INERTIAL COST

The constraint curvature analysis in Section II-B provides a clear view of the displacement produced by executing gait cycles, but does not contain any information about the resource or opportunity costs of executing the gaits. This cost information is critical for choosing optimal gaits. A gait that produces a large displacement per cycle but takes a large amount of energy or a very long time to execute is less useful than a gait that produces only a moderate displacement per cycle, but can be executed with a smaller energy expenditure or with higher frequency at the same instantaneous energetic cost, leading to a larger net displacement per energy or time.

Our systems’ inertial dynamics, described in Section III, suggest two basic means of measuring the effort required for system motion: the kinetic energy required to move with a given shape velocity, and the actuator force required to achieve a given acceleration. Integrating these quantities over one cycle of a gait produces an energetic cost that depends on the gait’s path through the shape space, its period, and its pacing.

For either effort cost, we can then assign an opportunity cost to each (path, pacing) pair, defined as the period T which, when combined with the path and pacing, produces a gait in which the time-averaged effort cost has unit value. As we discuss in Section V, this opportunity cost enables a well-posed gait optimization process, answering the question

“At a given level of effort, which path and pacing produces the fastest motion through the world?”

For the systems considered in this article, answering this question also answers the question

“For a given speed through the world, which combination of path, pacing, and period requires the least effort?”

but the former question allows for better decoupling between the “cost” and “benefit” of a gait (as discussed in Section V), and consequently simplifies reasoning about and finding optimal gaits.

Before directly considering the gait optimization problem, we find it useful to examine how the kinetic-energy and actuator-force measures of effort, under the constraint of unit-instantaneous effort, give rise to distinct “geometries of cost” in the resulting gait period. When combined with the constraint-curvature analysis from Section II-B, these geometries of cost provide a complete geometry of optimality for the systems’ gaits (discussed in Section V), in which the balance between

displacement-per-cycle and time-taken-per-cycle is encoded as a geometric balance between the constraint curvature enclosed by a gait and the length and curvature of the gait itself.

A. Kinetic-Energy Cost: Inertial Pathlength

If we take the system’s effort (instantaneous cost of motion) as its kinetic energy

$$E_{\text{instant}}^{\text{KE}} = \text{KE} \quad (29)$$

then we can make the following statements about the opportunity costs of gaits:

- 1) For any path and mean kinetic energy, the pacing that produces the shortest period T is the one in which the energy is held constant at this mean value. Because kinetic energy defines the metric-normalized speed, this means that the system moves at “constant speed” with respect to its inertial distribution.
- 2) Under the constant-inertial speed condition, the period T of a gait is proportional to its metric-weighted inertial pathlength and inversely proportional to the square root of the kinetic energy with which it is executed.

These properties are derived from the following geometric properties of the system dynamics:

- 1) *Optimal pacing*: Because of the quadratic relationship between \dot{r} and KE, trading off energy between portions of the gait (moving slower in one portion to move faster in another) at a given mean kinetic energy value always leads to a longer period T than is achieved by maintaining the kinetic energy constant at this mean value: the time gained in the faster sections is necessarily smaller than the time lost in the slower sections. A formal derivation of this principle is presented in Appendix A.
- 2) *Proportionality of period and pathlength*: If the energy in the system is held fixed at a given value, then the inertial pathlength S is equal to the product of the square root of kinetic energy and the gait period,

$$S = \int_0^T \sqrt{2\text{KE}} dt = \sqrt{\text{KE}} T \quad (30)$$

and the period is correspondingly equal to the inertial pathlength divided by the square root of the kinetic energy at which the gait is executed,

$$T = S/\sqrt{2\text{KE}}. \quad (31)$$

Fig. 6(a) illustrates the optimal pacing under the kinetic-energy measure of effort for a unit point mass moving around a “racetrack” path, which is to maintain constant speed through both straight and curved sections of a trajectory. Fig. 6(b) illustrates the linear proportionality between the radius of a circular path and the time a point mass requires to follow it at unit kinetic energy.

Fig. 7(a)–(d) illustrate how a nonisotropic inertia matrix affects the pacing in parameters of a constant-kinetic-energy trajectory. The rate at which gait period scales with respect to geometric scaling of the gait’s path (in shape-parameter space) depends both on the coordinate-stretch and on the underlying

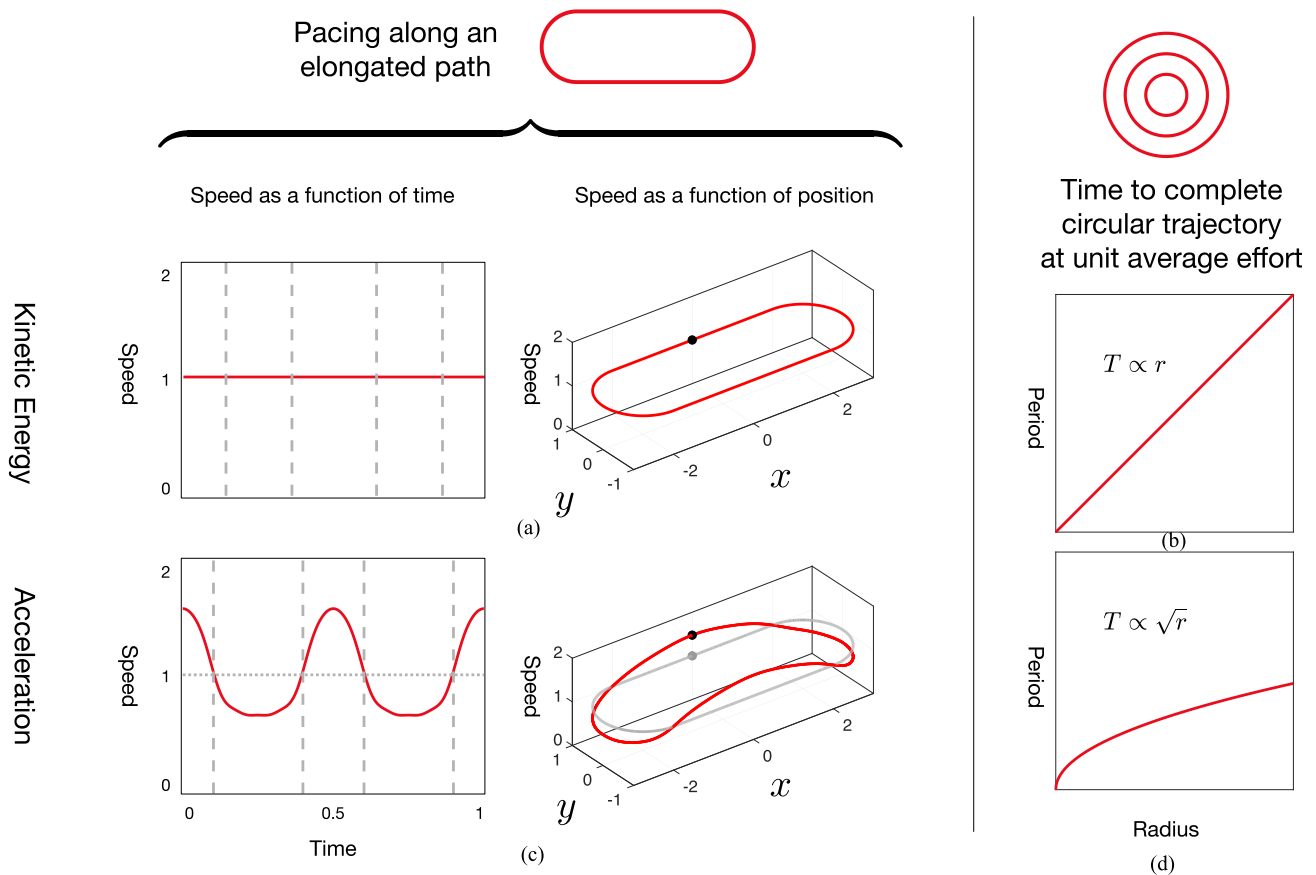


Fig. 6. Under a kinetic-energy measure of effort, (a) the optimal pacing for a point mass moving along an elongated path is to maintain constant kinetic energy, and (b) the time required for the point mass to move along a circular path at unit kinetic energy is proportional to the radius of the path. If we instead measure the effort required to follow the path as its squared acceleration, (c) the optimal pacing is slower in the curved sections than in the straight sections, and (d) the time required to follow the path at unit acceleration scales with the square root of the path radius.

(coordinate-invariant) curvature of the inertial manifold; as illustrated in Fig. 8(a)–(c), gaits for systems with positive Gaussian curvature (“domed” or “cupped” manifolds) experience slower growth of T as the amplitude of a cycle is scaled than do systems with flat inertial manifolds, whereas systems with negative Gaussian curvature (“saddle-shaped” manifolds) exhibit increased growth of T with respect to amplitude scaling.

Comparing the the system inertia matrix illustrated in Fig. 5 with the minimal-working-example in Fig. 7, our systems’ inertial geometries are stretched/compressed along the even axis, meaning that during a constant-energy trajectory, the joints move slowly as the system moves from \cap shapes to \cup shapes, and quickly as it moves from S shapes to Z shapes. The generally positive Gaussian curvature of the inertial geometry means that the pathlength, and, thus, period, of gaits can be expected to grow sublinearly with amplitude (which physically corresponds to the system being “curled more tightly” during higher amplitude gaits, and thus having a smaller overall moment of inertia).

B. Actuator-Force Cost: Covariant Acceleration

The kinetic-energy measure of effort measures the *net* work that the actuators must do on the system to move at a given speed, but does not account for the *individual* work that the actuators

perform as they shuffle kinetic energy between different moving pieces of the system. In most cases, this individual work cannot be regenerated or passed between actuators without significant losses, and in some cases actuators may be actively working against each other, with one actuator supplying energy to the system, while another dissipates energy.

To more accurately model the effort that the actuators put into the system, we can turn our attention to the actuator forces, which are related to the gait trajectories as in (15). A complete accounting of the actuator effort to produce those forces requires a detailed model of the actuators, but a good general-purpose model is to take the effort as the squared norm of the actuator forces,

$$E_{\text{instant}}^{\tau} = \tau^T \tau \quad (32)$$

capturing the idea that power consumption in an actuator (e.g., the resistive heat losses in an electric motor) grow superlinearly with respect to the force being supplied, and that the costs of producing forces in the individual actuators are decoupled.¹²

¹²If the cost of producing actuator force is not decoupled, or differs between actuators, this product could be further weighted to reflect such coupling or differing cost.

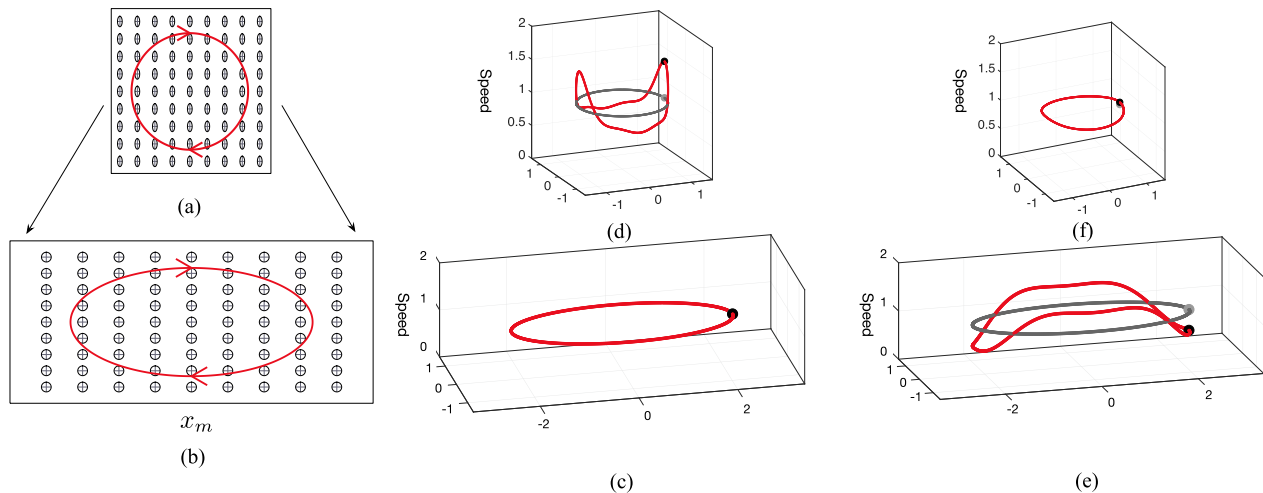


Fig. 7. If the inertia matrix is not isotropic (e.g., because the system has a shape-dependent distributed mass), then a trajectory that is circular in parameter space, as in (a), may not be truly circular as illustrated in (b). Consequently, the constant-kinetic-energy path in (c) appears to be variable-speed when viewed in coordinates (d) and the squared-covariant-acceleration-optimal pacing for the system (e) may appear to be constant speed when viewed in coordinates (f), even though the system actually experiences significant tangential acceleration.

As noted in (17), the actuator forces are the product of the system's inertia matrix and its covariant acceleration, such that the squared norm of actuator forces can be expressed in terms of the gait geometry (path and pacing) via the \mathcal{M}^2 norm of the covariant acceleration in (19). This proportionality of effort and acceleration results in a relationship between gait path, period, and pacing with the following properties:

- 1) For any path and mean squared covariant acceleration, the pacing that produces the smallest period is one in which the system moves slowly in highly curved sections and faster in straight sections (whereas under the kinetic-energy measure of effort, the optimal pacing has constant inertial speed).
- 2) For any path and mean squared torque, the pacing that produces the smallest period is biased from the acceleration-optimal pacing by a term that depends on the placement of the actuators on the system.
- 3) The period required to execute a gait with unit average effort is proportional to the fourth root of the effort required to execute the gait in unit time.
- 4) For geometrically similar gaits executed at optimal pacing, the period at unit-effort scales with the *square root* of the size of the gait (whereas period under the kinetic-energy measure of effort scales linearly with the size of the gait).
- 5) For gaits with the same inertial pathlength and mean squared torque, a “round” gait requires a shorter period than an “oblong” gait or one with “sharp corners” (whereas period under the kinetic-energy measure of effort is independent of the aspect ratio or distribution of curvature).

The properties are derived from the following geometric properties of the system dynamics:

- 1) *Acceleration-optimal pacing*: Using the metric-orthonormal representation of the acceleration norm from (26), the integral of squared covariant acceleration

can be written as

$$E_{\text{total}}^a = \int_0^T (\dot{v}^2 + (\kappa v^2)^2) dt. \quad (33)$$

The quadratic nature of this cost function means that as compared to moving with constant v , there necessarily exists a pacing that reduces the average acceleration by slowing down the system in sections where κ is large, even though this means accepting a nonzero contribution to acceleration from the \dot{v} term. A formal derivation of this principle is presented in Appendix B.

- 2) *Torque-optimal pacing*: In metric-orthonormal coordinates, the squared-torque cost is the integral of the torque-norm of acceleration from (28),

$$E_{\text{total}}^a = \int_0^T (b_{\dot{v}}\dot{v}^2 + 2b_{v\dot{v}}(\dot{v}\kappa v^2) + b_v(\kappa v^2)^2) dt. \quad (34)$$

This cost is weighted and biased relative to that in (33), but has qualitatively similar behavior: its quadratic nature means that relative to a constant-speed profile, shifting acceleration out of the $(\kappa v^2)^2$ and into a term with a \dot{v} term reduces the overall torque cost.

- 3) *Unit-effort/unit-time proportionality*: Changing the timescale on which the system motions occur by a factor c induces a change in the accelerations by a factor c^2 . This factor means that the normed actuator force τ_T during a gait with a given path and pacing and a period T is related to the actuator force during a gait with the same path and pacing but unit period, τ_1 , by a factor

$$\|\tau_T(t)\| = \frac{1}{T^2} \|\tau_1(t/T)\|. \quad (35)$$

Combining this relationship with a constraint that restricts our attention to gaits with unit average effort

$$\frac{\int_0^T \|\tau_T\|^2 dt}{T} = 1 \quad (36)$$

tells us that the opportunity cost (period) of a gait with a given path and pacing, executed with unit *average* effort satisfies the relationship

$$\frac{\int_0^T \left(\frac{1}{T^2} \|\tau_1(t/T)\|\right)^2 dt}{T} = 1. \quad (37)$$

Separating out the explicit factors of T in this expression provides an equation

$$T^5 = \int_0^T \|\tau_1(t/T)\|^2 dt \quad (38)$$

whose right-hand side can be rewritten to factor out T as

$$\int_0^T \|\tau_1(t/T)\|^2 dt = T \int_0^1 \|\tau_1(t)\|^2 dt \quad (39)$$

such that the period at unit average effort is revealed as being equal to the fourth root of the effort cost of the gait when it is executed with unit period

$$T = \left(\int_0^1 \|\tau_1\|^2 dt \right)^{1/4}. \quad (40)$$

- 4) *Square root proportionality of T and gait amplitude:* Increasing the pathlength of a gait by a factor ℓ , while maintaining its period, pacing, and curvature distribution scales $v(t)$ by a factor of ℓ , scales $\dot{v}(t)$ by a factor of ℓ^2 , and scales $\kappa(t)$ by $1/\ell$ such that the unit-time torque cost scales by ℓ^2 . Inserting this scaling factor into (40) then tells us that the unit-effort period for the gait increases by a factor of $(\ell^2)^{1/4} = \ell^{1/2}$.

Physically, this scaling factor corresponds to the property that as the size of the path increases, it takes proportionally more time to complete a cycle at a given v , but the path's curvature *decreases* as the gait grows. The system can thus move with a greater velocity while maintaining the same instantaneous κv^2 effort in the curves, so that the total time to complete the gait rises sublinearly.

- 5) *Optimality of roundness:* The quadratic nature of the cost function means that a gait with concentrated curvature (e.g., corners or an oblong shape) incurs greater $(\kappa v^2)^2$ costs at unit period, and, therefore, requires a greater period at unit effort, than does a system with the same inertial pathlength but more constant (rounder) curvature.

Fig. 6(c) illustrates how the factors described above influence the optimal pacing of a point mass moving around a “racetrack” path with unit mean squared acceleration, and Fig. 6(d) illustrates the square root proportionality between the radius of a circular path and the time the point mass requires to follow it at unit mean squared acceleration.

As illustrated in Fig. 7(e), the “stretch” from an anisotropic inertia matrix can lead to a “circular” gait in coordinates not being truly circular, such that its optimal pacing under the kinetic-energy metric proceeds around the circle at a nonconstant pace. Conversely, the stretch acts to mask the tangential acceleration

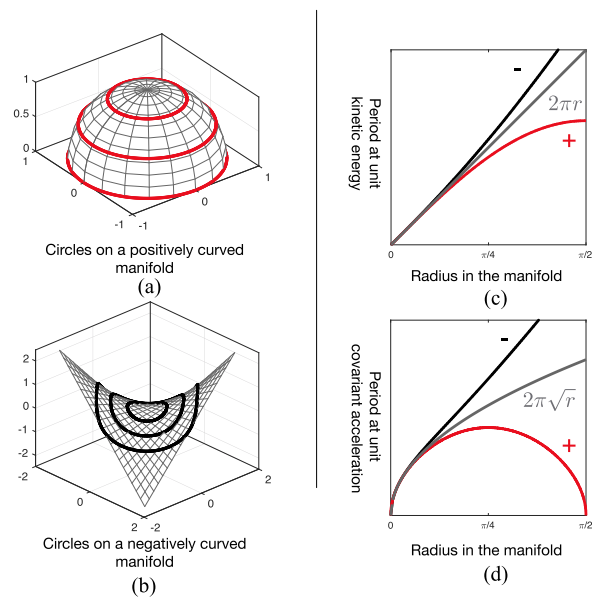


Fig. 8. Gaussian curvature of the inertial manifold affects the rate at which gait periods at unit effort grow with respect to the size of the gait. Positive curvature—the hemispherical geometry in (a)—slows the growth of the period with respect to gait size, and negative curvature—the saddle geometry in (c)—increases the rate of growth. These rates of growth are illustrated in (b) and (d), with the rate of growth in a “flat” space provided as reference.

Note that it is possible for a closed path to be a geodesic on the manifold—following the equatorial great circle on the hemisphere in (a)—in which case the period under the kinetic-energy measure of effort reaches a maximum (because it is proportional to inertial pathlength, and geodesics are curves at extrema of inertial pathlength). On such a path, the unit-time effort goes to zero (because geodesics are the trajectories of no acceleration, and the manifold's intrinsic curvature provides all of the “change of direction” required to close the loop) and the unit-effort period becomes undefined (because there is no periodic motion along a closed geodesic with nonzero constant effort).

(and, thus, effort) required to follow a circular path at mean squared acceleration, as is illustrated in Fig. 7(f): the stretch means that the acceleration-optimal pace requires slowing down in the “narrow ends” of the true elliptical gait geometry, but mapping the motion back into the parameter coordinates slows down the motion in the “straight sections” of the true geometry, such that the acceleration-optimal motion is constant-speed in the parameter space.

As illustrated in Fig. 8(d), the time required to follow a shape trajectory grows more slowly with loop size on inertial manifolds with positive (domed/cupped) curvature as compared to on a flat manifold, and grows more quickly on a manifold with negative (saddled) curvature. The reduction or increase in period for an acceleration-normalized trajectory is stronger than that for an energy-normalized trajectory; in the case of a manifold with sufficient positive curvature, the required period can actually shrink with increased loop size as the loops approach geodesics on the inertial manifold—the great circle on the hemisphere in Fig. 8(a), which requires zero force when followed at constant speed (and, therefore, does not have a well-defined “period at unit effort”).¹³

¹³This property is described by the Gauss–Bonnet theorem: The total curvature in a closed loop on a surface is equal to 2π minus the surface curvature enclosed by the loop.

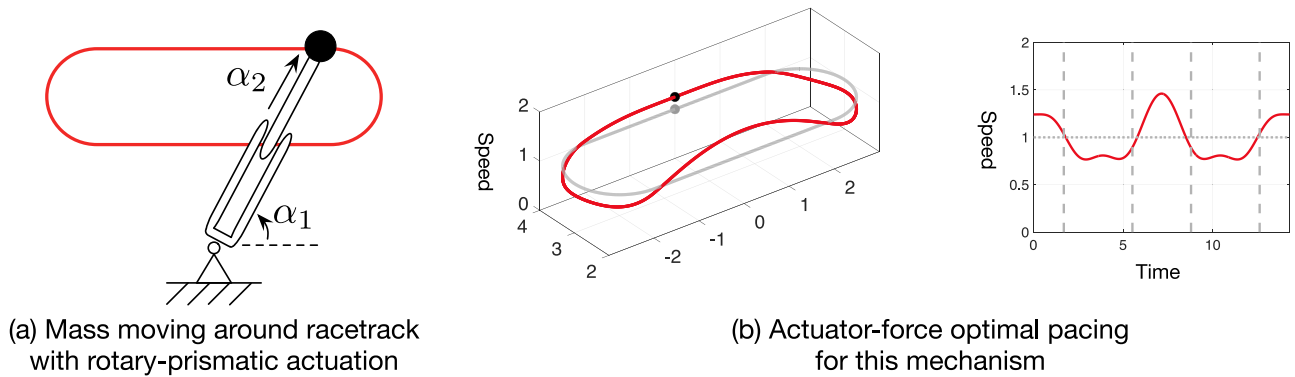


Fig. 9. Optimizing pacing with respect to squared actuator force is similar to optimizing with respect to covariant acceleration, but introduces a bias term based on the location of the actuators. (a) A point mass being moved around a race-track path by a rotary-prismatic mechanism. (b) The speed of the mass under actuator-force-optimal pacing as a function of both position along the racetrack, (left) and time within the trajectory (right). Note that the mass accelerates more aggressively in the straight section that is closer to the rotary joint than in the section that is further away, because linear acceleration along the track costs less torque when the moment arm is small. Similarly, the force-optimal pacing has a “speed wobble” during the curves, which correspond to accelerations along the racetrack becoming equivalent to accelerations directly along one of the mechanism joints.

The difference between optimal pacing under a raw covariant-acceleration cost and optimal pacing under an actuator-force cost can be understood by considering the case of a rotary-prismatic arm whose mass is concentrated at its distal end, as illustrated in Fig. 9. The point mass at the end of the arm has a flat inertial manifold, so its covariant acceleration is its simple acceleration on the path it follows, and this acceleration is equal (up to units) with the force that must be applied to the mass to produce the acceleration. When these forces are projected onto the joint angles, however, horizontal force requires increased α_1 torque when α_2 is large. This biasing term means that the optimal pacing involves a more aggressive acceleration during the “near” straight section, and also introduces a “speed wobble” to the curved sections based on how the mass’s acceleration becomes intermittently aligned with one or other of the mechanism joints.

Fig. 10 illustrates the effect of roundness on optimal gait speed. For paths of equal perimeter on a continuum between circles and squares, the optimal trajectory paces slow down at the corners (high curvature regions) to avoid incurring large centripetal acceleration costs. These decelerations mean that the system moves at a lower average speed than it moves through the circular path, thus requiring increased time to traverse nonround paths.

For the locomoting systems we are considering, the cancellation between true-geometry acceleration and coordinate-perceived speed from Fig. 7(e) and (f), means that the optimal pacing of a gait can be inferred directly from the “roundness” of its path in shape coordinates. The generally positive Gaussian curvature of the inertial surfaces means that the time taken to execute a gait at a given mean acceleration grows more slowly than the square root of the radius of the gait.

The torque bias appears as a metric-squared acceleration cost in the actuation coordinates, which for our systems means that accelerations along the even axis of the shape space become more costly, biasing the system toward constant-speed motion when moving along the even directions of the shape space and toward straighter motions along the odd axis of the shape space (so as not to have curvature in an even direction).

V. OPTIMAL GAITS

Combining the geometric relationship between a gait’s path and the displacement it produces (from Section II-B) with the geometric relationship between its path and pacing and its period at unit average effort (from Section IV) provides a geometric relationship between its path and pacing and its efficiency η , measured as its speed at unit effort,

$$\eta(\text{path}, \text{pacing}) = \frac{g_\phi(\text{path})}{T(\text{path}, \text{pacing})} \quad (41)$$

where, for conciseness of notation, we now use g_ϕ to indicate a signed norm of the displacement resulting from the gait (e.g., net rotation, or translation in a specific body direction).

The gradient of this efficiency with respect to a set of gait parameters p that define the path and pacing of the cycle is

$$\nabla_p \eta = \frac{1}{T} \nabla_p g_\phi - \frac{g_\phi}{T^2} \nabla_p T \quad (42)$$

which can be intuitively described as being the gradient of the net displacement with respect to the parameters, minus the gradient of the cost with respect to the parameters, with a normalizing factor to account for the different units of displacement and cost.

This gradient serves two purposes in our analysis. First, the gradients of g_ϕ and T can both be calculated (in closed or semiclosed form) in terms of the constraint curvature $\mathbf{D}(-\mathbf{A})$ and the inertia matrix M .¹⁴ Gradient-ascent optimization in the direction identified in (41) then allows for fast optimization of a high-density parameterization of the gait. Second, the gradient identified in (41) represents a fundamental truth about the system dynamics, capturing the *fundamental underlying structure* of any other optimization approach applied to these systems.

Under both the kinetic-energy and actuator-force measures of cost, the optimal gaits represent stable equilibria between the $\nabla_p g_\phi$ and $\nabla_p T$ terms. As we discussed in [9], these equilibria

¹⁴See [9] for calculation of the gradient of g_ϕ and the gradient of the inertial-pathlength instantiation of T , and Appendix C of this article for the gradient of the actuator-force instantiation of T .

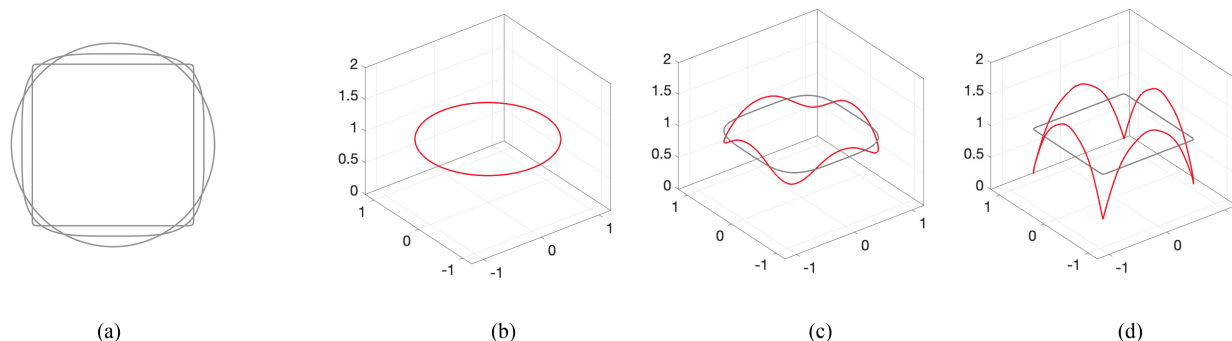


Fig. 10. For a set of paths with the same perimeter but different distributions of curvature (a), the optimal pacing has more variation in speed the sharper the corners in the trajectory are (b)–(d). These speed changes absorb some of the system’s acceleration budget, meaning that it must move slower (and, thus, take longer) as the curvature distribution along the path becomes more uneven.

(and, thus, the optimization process) resemble those defining the shape of soap bubbles—the $\nabla_p g_\phi$ term acts as an “inflating pressure” on the size of the gait, pushing it toward the maximum-displacement cycles (which follow zero-contours of the constraint curvature, fully enclosing sign-definite regions). The pathlength cost then acts as a “surface tension” (and for the force metric, a “bending strain”) term that pushes the gait to give up low-yield regions of the CCFs in favor of shorter (in the sense of both inertial pathlength and period) gaits that can be repeated more often around “richer” regions of the CCFs.

A. Optimal Gaits Under the Kinetic-Energy Metric

Under the kinetic-energy metric, the “surface tension” effect of the boundary cost acts as a true surface tension, penalizing gaits with a long inertial pathlength, but putting no direct penalty on curvature of the path. The way in which the kinetic-energy cost favors moving at constant inertial speed acts like the “concentration gradient” term in a soap bubble, which enforces an even distribution of material over the surface of the bubble. Optimal gaits under the kinetic-energy metric for the isolated three-link system and a set of perfect-fluid swimmers are illustrated in the first column of Fig. 11.

Note that for the piecewise and serpenoid swimmers, the optimal gaits cross slightly into positive regions of the constraint curvature at points along the odd shape axes. These excursions correspond to paths that are actually shorter on the manifold than the paths that exactly follow the zero contour in these regions; because the CCFs are small at these points, there is almost no displacement-per-cycle penalty for shortening the paths in this manner.

B. Optimal Gaits Under the Actuator-Force Metric

The optimization dynamics under the actuator-force measure of effort are similar to those under the kinetic-energy cost function, but with the following changes:

- 1) The “surface tension” is now based on the square root of the perimeter (rather than being directly proportional to it), so that the optimal gaits under the actuator-force metric are larger than those under the energy metric.

- 2) There is now a “bending stiffness” on the boundary in addition to the surface tension, such that optimal gaits under the actuator-force metric are rounder than those under the kinetic energy metric.
- 3) The pacing is now encouraged to slow down in more curved areas, rather than moving around the path at a constant inertial speed. Physically, this pacing change corresponds to constructing the gait curve from a material that contracts axially under applied bending load, producing a higher material density in curved regions.

Optimal gaits for the example systems under both the covariant-acceleration cost metric (which does not account for actuator placement) and the actuator-force metric (which does account for actuator placement) are illustrated in the second and third columns of Fig. 11.

Several of the acceleration- and force-optimal gaits cross out of sign-definite regions of the constraint curvature. Some of this behavior can be attributed to “cutting across” the sign-opposite region to reduce the pathlength of the gait (as in the kinetic-energy-optimal gaits); additionally, the “bending” component of the acceleration-based costs means that it may actually take more effort to avoid crossing into the sign-opposite region, such that the system is again encouraged to accept a penalty on displacement-per-cycle in order to achieve a better displacement-per-time for the gait.

C. Higher-Dimensional Shape Spaces

As we discuss in more detail in [9], the “soap bubble” descriptions of optimal gaits under the kinds of cost functions we discuss here extend directly to systems with more than two shape variables, with the key difference in higher dimensions being how we handle the curvature of the constraints.

In the case of system with three shape variables, the CCFs can be considered as 3-D vector fields, and the displacement of the system over a gait corresponds to the net flux of that field through the loop formed by the gait. The cost of executing a gait remains either its inertial pathlength or the squared norm of the actuator forces, and the optimal gaits are those that “catch the most flux” relative to the boundary cost.

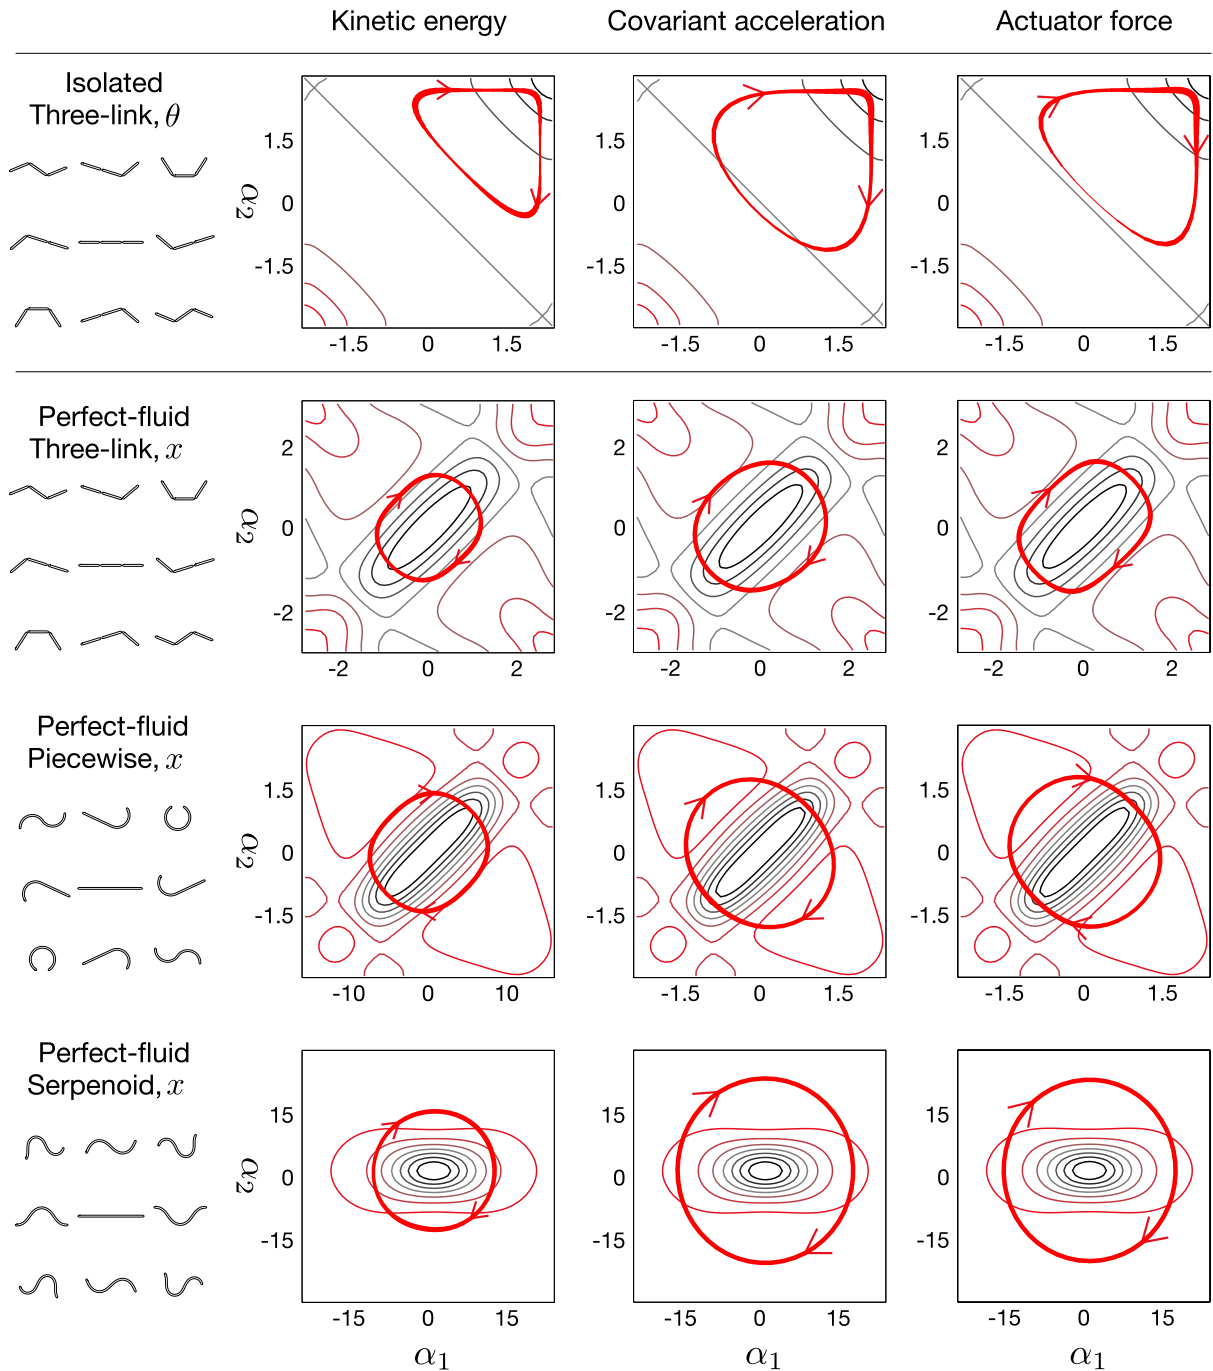


Fig. 11. Optimal gaits under the kinetic energy, covariant acceleration, and actuator-force metrics for a range of example systems. The optimal pacing is indicated via line thickness, with thick lines indicating “slow” changes in the shape variables and thin lines indicating “fast” changes (as if a set of points evenly distributed in time were “bunched up” or “stretched out”). The key trends to note are that the gaits optimized for kinetic energy are universally smaller and less rounded than the gaits optimized for covariant acceleration or actuator-force, and that the actuator-force gaits are shaped and paced so that they have large fast-moving segments in the “odd” dimension of the shape space (most visible on the isolated three-link and perfect-fluid piecewise gaits). This latter trend corresponds to the squaring of the inertia matrix in the actuator-force calculation, which reinforces the already existing property that motion along the odd axis is “easier” than motion along the even axis.

The structure of a 3-D flux field is difficult to visualize in a 2-D format, so, as discussed in [9], we find it convenient to represent the $\mathbf{D}(-\mathbf{A})$ fields via projection slices as in Fig. 12. For each system, we construct a surface that passes through the point, where the magnitude of the relevant (x, y, θ) component of $\mathbf{D}(-\mathbf{A})$ is largest, and aligned the surface so that it is normal

to that component of $\mathbf{D}(-\mathbf{A})$ at that point. Optimal gaits can be expected to lie close to this plane and encircle sign-definite regions, but bend out of the plane to fully capture the 3-D flux.

Optimal gaits for four-link (three joint) isolated and perfect-fluid systems are illustrated in Fig. 12(a) and (b). For the isolated system, the flux increases toward the corner of the shape space,




Fig. 12. Optimal gaits for systems with more than two shape variables. In the first two subfigures, the smaller gait is optimal with respect to kinetic energy and the larger gait is optimal with respect to actuator force; in the third subfigure, the two gaits follow similar enough paths to be visually indistinguishable. The colored planes are positioned such that they pass through the points where $\mathbf{D}(-\mathbf{A})$ is largest, and aligned so as to capture the most flux at those points. The colormaps are the CCFs indicating the flux of the relevant component of $\mathbf{D}(-\mathbf{A})$ through the planes. The optimal gaits for the systems are closely aligned with these planes and enclose rich, sign-definite areas within the planes (as if we had performed a 2-D optimization in the set of shape modes defined by the plane). The excursions of the optimal gaits from the planes are due to two causes: The cost metrics on the system favor motions in certain planes, and $\mathbf{D}(-\mathbf{A})$, as a 3-D or 4-D differential form, changes direction and value across the space and is not completely characterized by the illustrated plane. For the four-mode serpenoid system, the motion of the fourth shape variable is indicated via line thickness.

and the resulting gaits thus force themselves into the corner as much as possible (similarly to the isolated-system gaits in Fig. 11). On the perfect-fluid system, the optimal gaits lie very close to the projection plane, but curl slightly outward. This divergence from the center-point-maximum-flux plane both reduces the cost of the motion (because the plane contains high-cost even-mode motions), and allows the gaits to better capture flux away from the center of the shape space (because the CCF can change direction across the space).

Beyond three dimensions, the CCFs cannot be treated as vectors, but instead must be treated in a fully differential-geometric fashion as “differential two-forms.” At a practical level, this just means that the flux of $\mathbf{D}(-\mathbf{A})$ through the gait loop is directly associated with a surface bounded by the loop, rather than a vector field passing through such a surface,¹⁵ and the geometry of optimal gaits remains the same: they capture the most $\mathbf{D}(-\mathbf{A})$ flux for the least boundary cost. Optimal gaits for a perfect-fluid system with four serpenoid modes are illustrated in Fig. 12(c).

D. System Comparison

As illustrated in Fig. 13(a), efficiency under the kinetic-energy cost function increases as the swimmer becomes more continuous (gains extra deformation modes or adopts smoother deformation modes), following the same pattern we observed for the viscous swimmers in [9]. The swimmer with discrete joints gains the most efficiency (because the gain in continuity from adding a degree of freedom is relatively large), whereas the serpenoid swimmer barely gains any efficiency (because it is already continuous, and the extra shape modes do not add significantly to its swimming capability). The piecewise-continuous system falls between the other two (as should be

expected), and at high degrees of freedom converges on the same efficiency as achieved by the serpenoid systems.

As illustrated in Fig. 13(b) efficiency under the covariant-acceleration cost function exhibits a different pattern: for systems that are relatively discontinuous (discrete-joint and low-degree-of-freedom piecewise systems), adding a degree of freedom improves efficiency, but for systems that are relatively continuous (serpenoid and high-degree-of-freedom piecewise systems), adding a DoF actually reduces the efficiency.

Referring to the optimal gaits plotted in Fig. 12(c), we can explain this behavior by noting that the optimal gaits lie almost in the plane of the original modes, such that there are no strong kinematic benefits from moving with the additional modes. For the kinetic-energy metric, the presence of the extra modes does not incur any additional costs (because the optimal gait simply does not move along these directions). For the covariant-acceleration metric, however, the additional modes do incur additional cost: The system must use part of its force budget to actively brace against the natural accelerations in those directions, reducing the acceleration allocated to each DoF.

This reduced force allocation directly means that the system must move more slowly through its shape space to satisfy the unit-average-acceleration condition (meaning that adding a DoF in these situations moves the system to the right in the displacement/time plot). The reduced force allocation also indirectly encourages the system to adopt gaits that avoid generating forces that project into the new DoFs, trading off performance in favor of lower costs (and moving down on the displacement/time plot).

We have not plotted relative efficiency under the torque-based cost, because a fair comparison would require construction-specific details about the system (as opposed to the two costs that we do plot, which depend only on the physical geometry of the system and are independent of the actuator placement and internal workings).

¹⁵See [9] for further discussion of this point.

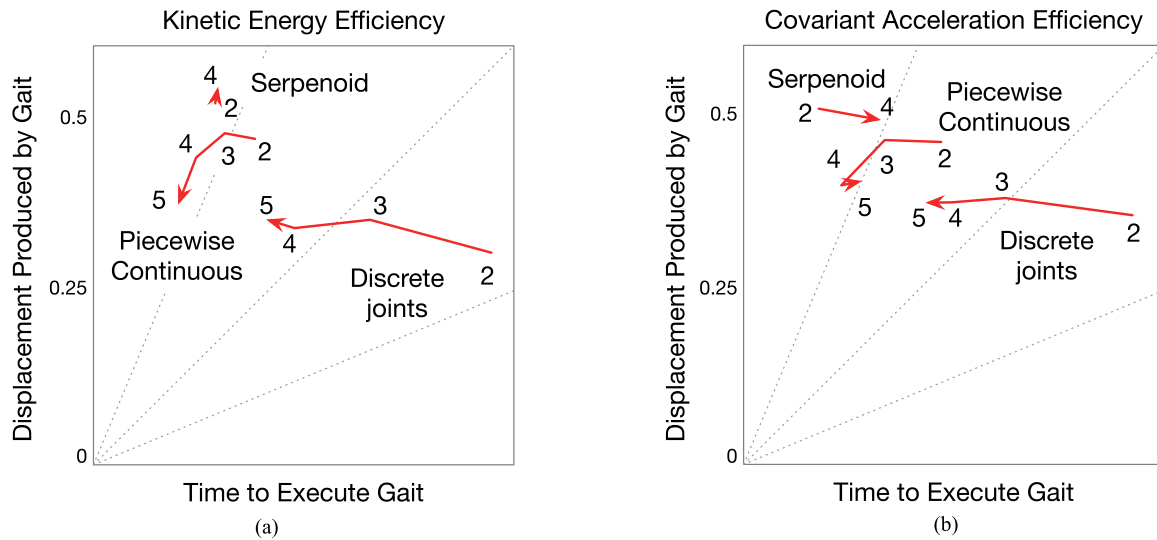


Fig. 13. Comparison of the most efficient x -translation gaits for discrete-joint, piecewise-continuous, and serpenoid swimmers, measured as the displacement they produce (in body lengths) divided by the time required to execute them under the constraints of unit kinetic energy and unit average covariant acceleration. The steeper the line from the origin to the system, the more efficient its optimal gait is. Under the kinetic-energy cost function, efficiency increases with both mode count and continuity of the modes. Under the covariant-acceleration cost function, efficiency increases with respect to continuity of the modes, but decreases with respect to mode count for systems that are already highly continuous (because the new modes do not significantly increase the capability of the system, but it takes effort to avoid moving along them). Note that the vertical scales on the two plots are comparable, but the horizontal scales are not (because there is no inherent relationship between “unit kinetic energy” and “unit average covariant acceleration”). The dotted lines in the plots serve as grid-lines or references; all points on any of these lines are equivalent in terms of displacement-per-time at unit-average-cost, and gaits above/left of a given line are more efficient than gaits that are below/right of that line.

VI. DIRECT NONHOLONOMIC CONSTRAINTS

In the text above, we extracted both the locomotion model (encoded in \mathbf{A}) and the cost model (encoded in \mathcal{M}_r) from the kinetic energy metric \mathcal{M} over the full (shape and position) configuration space. There also exist systems for which the inertial costs of motion we present in this article make good cost models, but whose locomotion models are not generated by applying conservation of momentum to the system. A canonical example of such systems is the three-link *kinematic snake* [13], [15], [17], [20], [27], which, as illustrated in Fig. 14(a), is a three-link chain with a passive wheelset attached to each link. These wheels act as direct nonholonomic constraints on the system motion (prohibiting system velocities in which any wheelset slides laterally) and together these constraints define a local connection \mathbf{A} for the system of the same form as in (8), whose structure is illustrated in Fig. 14(b).

Although this local connection is not itself derived from the system’s inertial dynamics, we can use it to pull back the full inertia matrix into a reduced inertia matrix via the same process as in (13). This reduced inertia matrix then provides a metric for evaluating the kinetic energy that it takes to move with a given shape velocity in the presence of the wheel constraints, and how much covariant acceleration and actuator force is required to follow a given shape trajectory.

The constraint curvature for the kinematic snake can be analyzed just as it was for the other two-joint systems, except on the $\alpha_1 = \alpha_2$ line, where the local connection becomes singular, as illustrated in Fig. 14(c).¹⁶ Physically, these singularities

¹⁶Singularities also appear at $\alpha_1 = \pm\pi$ and $\alpha_2 = \pm\pi$, but these self-colliding configurations are outside the scope of our discussion here.

correspond to the system forming a shape in which the lines normal to the wheel constraints meet at a single point, as illustrated in Fig. 14(d). This alignment means that the system can rotate around the convergence point (move in position space without changing shape) and can only cross the $\alpha_1 = \alpha_2$ line at right angles to it.

Mathematically, the singularity appears in the system dynamics as the local connection \mathbf{A} becoming asymptotically parallel to the singularity, with magnitude increasing toward infinity at the singularity, and with opposite sign on either side of the singularity. In the constraint curvature $\mathbf{D}(-\mathbf{A})$, which is calculated as a derivative of $-\mathbf{A}$, the singularity appears in the x component as an asymptotic rise in magnitude with the *same* sign on both sides of the singularity, but with an opposite-sign “delta function” at the singularity.

This delta-function structure corresponds to the “well” at the center of the perfect-fluid system’s $\mathbf{D}(-\mathbf{A})^x$ CCF in Fig. 4. As the ratio of lateral added mass to longitudinal and rotational added mass increases, this well becomes increasingly narrow; the kinematic snake represents a limit-case in which the lateral added mass becomes infinitely large, providing infinite resistance to the lateral motion of the links.

The integral of the constraint curvature inside a gait for the kinematic snake includes both the surface integral used in the earlier examples in this article and a line integral of delta functions over the segment of the singularity that lies within the gait. Including this segment (which encodes the discrete change in \mathbf{A} across the singularity) in the integral is analogous to treating corners in a line or creases in a surface as discrete equivalents of their geometric curvature when integrating to find net change in angle along a line or the net solid angle subtended by a surface.

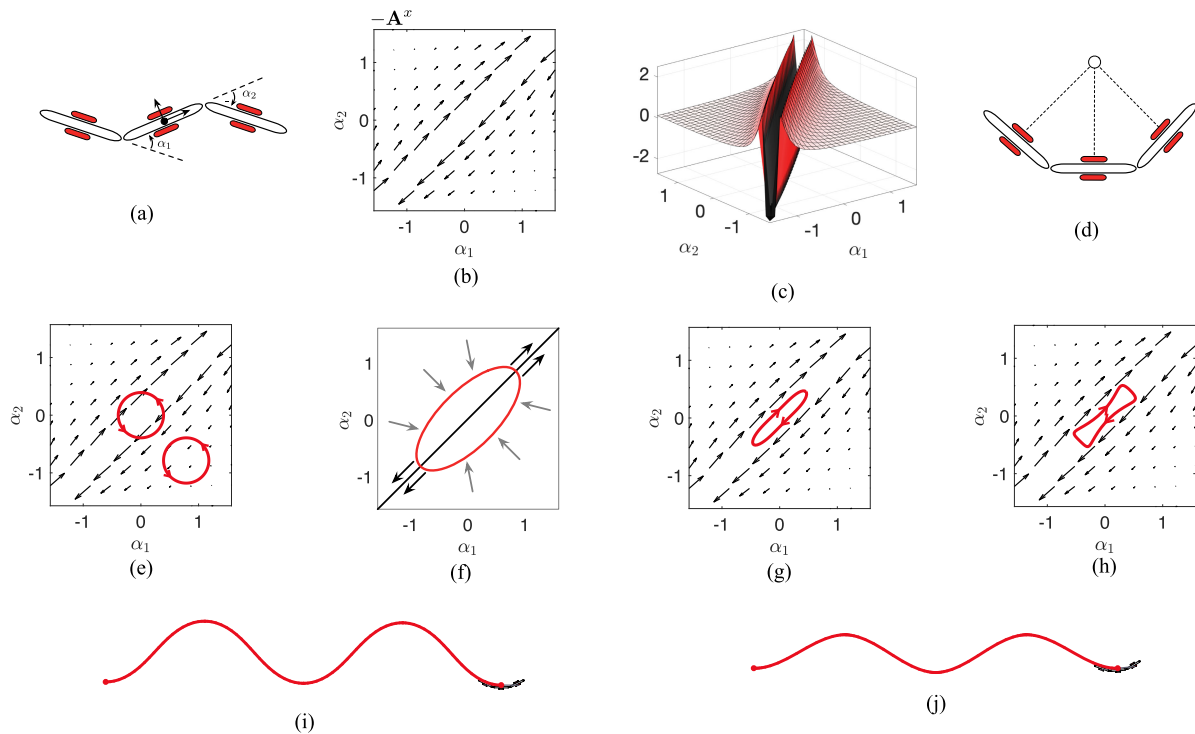


Fig. 14. (a) Three-link kinematic snake has a set of passive wheels on each link, preventing the link from moving laterally. (b) The resulting local connection and (c) curvature of the local connection have discontinuities at the $\alpha_1 = \alpha_2$ line, corresponding to (d) configurations in which the lines of constraint pass through a common center. Gaits can only cross this singularity orthogonal to it; as illustrated in (e), positively oriented gaits that cross the singularity go “against the flow” of the field, whereas gaits that do not cross the singularity go “with the flow.” For singularity-crossing gaits, the optimization pressure from the CCF is to stretch along the discontinuity (enclosing the singular “valley” in the CCF), while pinching inward everywhere else (excluding the opposite-sign regions on either side of the valley), as illustrated in (f). The specific shapes of the optimal gaits under kinetic-energy (g) and actuator-force (h) measures of effort are driven by their “membrane-like” and “beam-like” dynamics; in particular the gait in (h) pulls away from the singularity to avoid the high cost of accelerating in its vicinity. As illustrated in (i) and (j), the optimal gaits trace out sinuous paths through the world.

Note that as illustrated in Fig. 14(e), the net line integral on $-\mathbf{A}$ for a gait that crosses the singularity is negative for counterclockwise loops (and corresponds to the change in direction of the field across the singularity), whereas the integral for noncrossing gaits is positive for counterclockwise loops (and corresponds to the fields’ change in magnitude within each region). This difference highlights the fundamental importance of including the singularity’s effect on $\mathbf{D}(-\mathbf{A})$ in any constraint curvature analysis involving gaits that cross the $\alpha_1 = \alpha_2$ line. In particular, it is important to recognize that although the limit of $\mathbf{D}(-\mathbf{A})^x$ is *positive* as the singularity is approached from both sides, the singularity itself is sufficiently *negative* so as to make any integral containing it also negative; simply “bridging over” the singularity by taking the value of $\mathbf{D}(-\mathbf{A})$ as the mean of its limit value on either side would result in an incorrect prediction of the direction of net displacement.

As illustrated in Fig. 14(f), the line integral applies pressure on gait curves to expand along the $\alpha_1 = \alpha_2$ line, whereas the opposite-sign constraint curvature elsewhere in the shape space pushes the gait inwards. The optimal gaits for this system under the kinetic-energy and actuator-force metrics illustrate the effects of these pressures. The gait curves stay very close to the singularity, minimizing the amount of positive constraint curvature they enclose, which would otherwise cancel out some of the benefit of enclosing the singularity (On the plot of $-\mathbf{A}^x$,

we see that staying close to the singularity means the gait is staying in a region where the magnitude of $-\mathbf{A}^x$ is large).

The optimal gait under the kinetic-energy measure of effort, illustrated in Fig. 14(g), has an oblong shape, and is prevented from straying too close to the singularity by the large amount of kinetic energy that changing shape in near-singular configurations entails. The optimal gait under the force-based measure of energy, illustrated in Fig. 14(g), bulges out at the ends to reduce the system’s covariant acceleration, while changing direction. The growth along the singularity of both of these gaits is ultimately constrained by the need to prevent their world trajectories, illustrated in Fig. 14(i) and (j) from becoming too sinuous, wasting motion in the lateral directions.

VII. CONCLUSION

In this article, we have presented a set of geometric principles defining the shape of optimal trajectories for isolated and perfect-fluid locomoting systems under “least-action,” “least-squared-acceleration,” and “least actuator effort” objective functions.

A key feature of this geometric framework is that it provides a fair comparison between “small” gaits that can be repeated at high frequency and “large” gaits that produce more displacement, but also require more time to execute each cycle. Fixing the average instantaneous cost of motion assigns each gait path a

best-case-time-to-execute “opportunity cost” related to its size; dividing the gait’s induced displacement by this time provides a measure of gait effectiveness that is independent of any artifacts of a “displacement per cycle” analysis.

In comparing the covariant acceleration cost function with the actuator force cost function, this framework also highlights the difference between optimizing for properties that are fundamental properties of the mechanism and those that also depend on how the actuators are attached to the mechanism.

Our extension of the inertial cost functions to nonholonomically constrained systems, such as the kinematic snake highlights the independence of the inertial cost formulation from the inertial dynamics formulation. In addition, it demonstrates that the cost formulations continue to provide geometric insight in the presence of singularities in the dynamics.

Together with our previous development of such rules for systems with viscosity-dominated physics, we feel that this work provides a “complete picture” of the planar locomotion of these kinematic locomoting systems, relating the path and pacing of their optimal gaits to the fundamental geometric structure of their physics.

Although we see this work as “completing” one line of research, we do not see it as “closing off” work in this area. Rather, we see it as establishing a “well-furnished basecamp” from which to stage further investigations.

Potential avenues for further work (some of which we and others are currently investigating) include using the frameworks discussed in this article to understand and analyze observed motions of physical systems, analysis of hybrid systems that can make discrete changes in their contact state with the environment, and the development of steering and other task-focused locomotion algorithms that apply feedback and planning rules on top of the motion primitives identified in our gait analysis.

APPENDIX A ENERGY-OPTIMAL PACING

In Section IV-A, we use the quadratic nature of the kinetic-energy cost function to justify “constant inertial speed” as a property of trajectories that minimize time to traverse a path at a given average kinetic energy. Taking the period of motion as T and its inertial pathlength as S , the corresponding optimization can be stated formally as

$$\text{minimize } T = S/\bar{v} \quad (43)$$

$$\text{subject to constraint } \frac{1}{T} \int_0^T v^2(t) dt = K \quad (44)$$

where v is the system’s inertial speed, \bar{v} is its time-averaged value, and K is a constant. Expanding this constraint in terms of a fluctuating velocity

$$v(t) = \bar{v} + \tilde{v}(t) \quad (45)$$

with the speed fluctuation constrained as

$$\int_0^T \tilde{v}(t) dt = 0 \quad (46)$$

gives

$$\frac{1}{T} \bar{v}^2 T + 0 + \int_0^T \tilde{v}^2(t) dt = K \quad (47)$$

and thus,

$$\bar{v}^2 = K - \int_0^T \tilde{v}^2(t) dt. \quad (48)$$

Because all terms in (47) are nonnegative, K is necessarily greater than or equal to either of the other terms. The squared mean speed is thus maximized [and hence the time-to-execute cost in (43) is minimized] when $\tilde{v}(t) = 0$ such that the system moves with constant inertial speed.

APPENDIX B ACCELERATION-OPTIMAL PACING

In Section IV-B, we use the quadratic nature of the kinetic-energy cost function to justify “slowing down in curved sections” as a property of trajectories that minimize time to traverse a path at a given average covariant acceleration.

To formally demonstrate this property, we can first use the unit-time/unit effort quartic proportionality from (40) to convert the constraint on the gaits we consider from fixed-effort to fixed-time. We can then convert the cost function from an integral over time to an integral over the path as

$$E^a = \int_0^T (\dot{v}^2(t) + \kappa^2(t)v^4(t)) dt \quad (49)$$

$$= \int_0^S \frac{\partial t}{\partial s} (\dot{v}^2(s) + \kappa^2(s)v^4(s)) ds \quad (50)$$

$$= \int_0^S ((v')^2 v(s) + \kappa^2(s)v^3(s)) ds \quad (51)$$

where the final step is based on the identities $\frac{\partial t}{\partial s} = 1/v(s)$ and $\dot{v}(s) = \frac{\partial v(s)}{\partial s} v(s) = v'(s)v(s)$.

The derivative of this cost with respect to a variation δ in the velocity $v(s)$ is

$$\frac{\partial E^a}{\partial \delta} = \int_0^S 2v' \left(\frac{\partial}{\partial \delta} v' \right) + (v')^2 \left(\frac{\partial}{\partial \delta} v \right) + \kappa^2(3v^2) \left(\frac{\partial}{\partial \delta} v \right) ds \quad (52)$$

$$= \int_0^S 2v' \left(\frac{\partial}{\partial \delta} v' \right) + ((v')^2 + (3\kappa^2 v^2)) \left(\frac{\partial}{\partial \delta} v \right) ds. \quad (53)$$

We demonstrate that moving with constant inertial speed is not an equilibrium trajectory for acceleration-cost with respect to velocity variations that move the system slower in curved portions of the path by noting that for a constant-speed trajectory $v(s) = v_0$, the accompanying $v'(s) = 0$ condition means that the derivative in (53) does not depend on $(\frac{\partial}{\partial \delta} v')$

$$\left. \frac{\partial E^a}{\partial \delta} \right|_{v(s)=v_0} = 3v_0^2 \int_0^S \kappa^2(s) \left(\frac{\partial}{\partial \delta} v(s) \right) ds \quad (54)$$

meaning that there is locally no cost for variations introducing v' speed-spatial acceleration, but that there is a benefit to decreasing speed in more-curved sections of the path. To see this

benefit directly, we can introduce a velocity variation

$$v(s, \delta) = (1 - \delta \kappa^2(s)) \left(\frac{1}{T} \int_0^S \frac{1}{1 - \delta \kappa^2(s)} ds \right) \quad (55)$$

in which the left-hand term reduces the speed in proportion to the squared curvature of the path, and the right-hand term rescales the resulting velocities to keep the mean velocity constant, and which has $v'(s) = 0$ when $\delta = 0$. The derivative of the velocity with respect to the variations at $\delta = 0$ resolves to

$$\left. \frac{\partial v(s, \delta)}{\partial \delta} \right|_{\delta=0} = v_0 (\kappa_{\text{av}}^2 - \kappa^2(s)) \quad (56)$$

(where $v_0 = S/T$), i.e., the system slows down in regions that have a greater-than-average squared curvature, while speeding up in regions with smaller-than-average squared curvature. Inserting this velocity derivative into the cost derivative from (54) gives the cost derivative with respect to this variation as

$$\left. \frac{\partial E^a}{\partial \delta} \right|_{\delta=0} = 3v_0^2 \int_0^S \kappa^2(s) (\kappa_{\text{av}}^2 - \kappa^2(s)) ds. \quad (57)$$

Because the $\kappa^2(s)$ term scales negative regions of $(\kappa_{\text{av}}^2 - \kappa^2(s))$ by a greater magnitude than it scales positive regions, the integral is guaranteed to be nonpositive, and is strictly negative outside of the constant-squared-curvature case, where $\kappa^2(s) = \kappa_{\text{av}}^2$.

APPENDIX C

GRADIENT OF THE TIME PERIOD FOR ACTUATOR-FORCE COST

The second term of (42) takes $\nabla_p T$ as a measure of how moving points in the gait path gait influences the acceleration cost of executing the gait. Taking the gradient of the actuator-force-constrained period T from (40) with respect to the parametrization p gives the gradient of the gait's period in terms of the unit-period cost of the gait at with the current parameters and the gradient of the unit period cost with respect to the parameters

$$\nabla_p T = \nabla_p \left(\int_0^1 \tau_1^2 dt \right)^{1/4} \quad (58)$$

$$= \frac{1}{4} \left(\int_0^1 \tau_1^2 dt \right)^{-3/4} \left(\nabla_p \left(\int_0^1 \tau_1^2 dt \right) \right) \quad (59)$$

which, by replacing the first integral expression with T via (40), simplifies as

$$\nabla_p T = \frac{1}{4T^3} \int_0^1 (2\tau_1 \nabla_p \tau_1) dt. \quad (60)$$

Because the quantities T and τ_1 are known at each step in the optimization, computing the gradient $\nabla_p T$ requires only further calculation of $\nabla_p \tau_1$. Applying the gradient ∇_p to each of the terms in (15) and (16) via the chain rule gives

$$\nabla_p \tau_1 = (\nabla_p \mathcal{M}_r(r)) \ddot{r} \quad (61a)$$

$$+ \mathcal{M}_r(r) (\nabla_p \ddot{r}) \quad (61b)$$

$$+ \sum_i \left(\left(\nabla_p \frac{\partial \mathcal{M}_r(r)}{\partial r_i} \right) \dot{r}_i + \frac{\partial \mathcal{M}_r(r)}{\partial r_i} (\nabla_p \dot{r}_i) \right) \dot{r} \quad (61c)$$

$$+ \sum_i \left(\frac{\partial \mathcal{M}_r(r)}{\partial r_i} \dot{r}_i \right) (\nabla_p \dot{r}) \quad (61d)$$

$$- \frac{1}{2} \left((\nabla_p \dot{r})^T \frac{\partial \mathcal{M}_r(r)}{\partial r} \dot{r} + \dot{r}^T \left(\nabla_p \frac{\partial \mathcal{M}_r(r)}{\partial r} \right) \dot{r} \right) \quad (61e)$$

$$- \frac{1}{2} \dot{r}^T \frac{\partial \mathcal{M}_r(r)}{\partial r} (\nabla_p \dot{r}) \quad (61f)$$

from which $\nabla_p \tau_1$ can be obtained by constraining the gait defined by r to be completed in unit time.

This expression may be further rearranged such that the gradient operation with respect to the parameters is applied only to the path variables by expanding $\nabla_p \mathcal{M}_r(r)$ and $\nabla_p \frac{\partial \mathcal{M}_r(r)}{\partial r}$ using the chain rule as

$$\nabla_p \mathcal{M}_r(r) = \left(\frac{\partial \mathcal{M}_r(r)}{\partial r} \right) (\nabla_p r) \quad (62)$$

$$= \sum_i \left(\frac{\partial \mathcal{M}_r(r)}{\partial r_i} \nabla_p r_i \right) \quad (63)$$

and

$$\nabla_p \frac{\partial \mathcal{M}_r(r)}{\partial r} = \left(\frac{\partial^2 \mathcal{M}_r(r)}{\partial r^2} \right) (\nabla_p r) \quad (64)$$

$$= \begin{bmatrix} \sum_j \left(\frac{\partial^2 \mathcal{M}_r(r)}{\partial r_1 \partial r_j} \nabla_p r_j \right) \\ \vdots \\ \sum_j \left(\frac{\partial^2 \mathcal{M}_r(r)}{\partial r_n \partial r_j} \nabla_p r_j \right) \end{bmatrix}. \quad (65)$$

APPENDIX D

ACTUATOR-FORCE METRIC FOR VISCOUS SYSTEMS

In [9] and [16], we measured the cost of a gait for a viscous system as the energy dissipated through viscous friction between the swimmer and its environment. This cost translated into geometric form as the pathlength of the gait under a Riemannian metric M_r constructed by reducing the total system drag matrix via the local connection. This drag cost is analogous to the covariant-acceleration cost for the inertial systems, in that it considers only the force acting on each particle of the system at that particle's location, but does not account for the leverage the actuators have on the body. The actuator-force cost for the viscous system is the pathlength under the *square* of the reduced drag matrix,

$$\text{cost} = \int_{\phi} \overbrace{\dot{r}^T M}^{\tau^T} \overbrace{M \dot{r}}^{\tau}. \quad (66)$$

The optimal gaits under this squared metric are qualitatively the same as those under the original metric, and are slightly shorter (when viewed in the parameter space) along the even axes of the shape space (where the true metric geometry is

“long” relative to the coordinate geometry). The efficiency under this cost function follows the same pattern as the covariant acceleration efficiency in Fig. 13(b). Increasing mode-count for low-continuity systems increases their actuator-force efficiency, but adding modes to systems that are already highly continuous reduces their efficiency because the system must actively use force to stop unused modes from moving.

REFERENCES

- [1] R. Montgomery, “Gauge theory of the falling cat,” in *Dyn. Control Mech. Systems: Falling Cat Related Problems*, M. J. Enos, Ed., Providence, RI, USA: Amer. Math. Soc., 1993.
- [2] G. C. Walsh and S. Sastry, “On reorienting linked rigid bodies using internal motions,” *IEEE Trans. Robot. Automat.*, vol. 11, no. 1, pp. 139–146, Feb. 1995.
- [3] E. Kanso, J. E. Marsden, C. W. Rowley, and J. B. Melli-Huber, “Locomotion of articulated bodies in a perfect fluid,” *J. Nonlinear Sci.*, vol. 15, pp. 255–289, 2005.
- [4] J. B. Melli, C. W. Rowley, and D. S. Rufat, “Motion planning for an articulated body in a perfect planar fluid,” *SIAM J. Appl. Dynam. Syst.*, vol. 5, no. 4, pp. 650–669, Nov. 2006.
- [5] 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.
- [6] O. Itani and E. Shamma, “Motion planning for a redundant planar snake robot,” in *Proc. IEEE/ASME Int. Conf. Adv. Intell. Mechatron.*, 2020, pp. 1483–1488.
- [7] S. Ramasamy and R. L. Hatton, “Soap-bubble optimization of gaits,” in *Proc. IEEE Conf. Decis. Control*, Las Vegas, NV, USA, 2016, pp. 1056–1062.
- [8] R. L. Hatton, B. Bittner, S. Ramasamy, and S. Revzen, “Data-driven geometric gait analysis,” in *Proc. Dyn. Walking*, Jun. 2017, pp. 1–2.
- [9] S. Ramasamy and R. L. Hatton, “The geometry of optimal gaits for drag-dominated kinematic systems,” *IEEE Trans. Robot.*, vol. 35, no. 4, pp. 1014–1033, Aug. 2019.
- [10] A. Shapere and F. Wilczek, “Gauge kinematics of deformable bodies,” *Amer. J. Phys.*, vol. 57, no. 6, pp. 514–518, Jun. 1989.
- [11] R. M. Murray and S. S. Sastry, “Nonholonomic motion planning: Steering using sinusoids,” *IEEE Trans. Autom. Control*, vol. 38, no. 5, pp. 700–716, May 1993.
- [12] S. D. K. Kelly and R. M. Murray, “Geometric phases and robotic locomotion,” *J. Robot. Syst.*, vol. 12, no. 6, pp. 417–431, 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] 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, 2008, Art. no. 063016.
- [15] 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.
- [16] R. L. Hatton, T. Dear, and H. Choset, “Kinematic cartography and the efficiency of viscous swimming,” *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 523–535, Jun. 2017.
- [17] R. L. Hatton and H. Choset, “Nonconservativity and noncommutativity in locomotion,” *Eur. Phys. J. Special Topics: Dyn. Animal Syst.*, vol. 224, no. 17–18, pp. 3141–3174, 2015.
- [18] 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.
- [19] 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, Dec. 2007.
- [20] E. A. Shamma, 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.
- [21] 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.*, 1998, pp. 1–4.
- [22] W. Magnus, “On the exponential solution of differential equations for a linear operator,” *Commun. Pure Appl. Math.*, vol. 7, pp. 649–673, 1954.
- [23] R. Montgomery, *A Tour of Subriemannian Geometries, Their Geodesics, and Applications*. Providence, RI, USA: American Mathematical Society, 2002.
- [24] J. B. Tenenbaum, V. De Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [25] M. Tissot, *Memoire Sur La Representation Des Surfaces Et Les Projections Des Cartes Geographiques*. Paris, France: Gautier-Villars, 1881.
- [26] E. Abbena, S. Salamon, and A. Gray, *Modern Differential Geometry of Curves and Surfaces with Mathematica, Series Textbooks in Mathematics*. Boca Raton, FL, USA: CRC, 2017.
- [27] T. Dear, B. Buchanan, R. Abajian-Guerrero, S. D. Kelly, M. Travers, and H. Choset, “Locomotion of a multi-link non-holonomic snake robot with passive joints,” *Int. J. Robot. Res.*, vol. 39, no. 5, pp. 598–616, 2020.




Ross L. Hatton received the Ph.D. and M.S. degrees in mechanical engineering from Carnegie Mellon University, Pittsburgh, PA, USA, and the SB degree in mechanical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA.

He is currently 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.



Zachary Brock received the B.S. degree in mechanical engineering from Brigham Young University, Provo, UT, USA, and the M.S. degree in robotics from Oregon State University, Corvallis, OR, USA.

He has been with Daimler Trucks North America, Portland, OR, USA, as an Autonomous Engineer and currently holds a position with Torc Robotics, Inc., Blacksburg, VA, USA, as a Software Engineer working to develop fully autonomous semitrucks.



Shuoqi Chen received the M.S. degree in robotics from Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.

He is a Navigation Engineer with Intuitive Surgical, Inc., Sunnyvale, CA, USA. His research interests include locomotion and geometric motion planning of multilink robots.



Howie Choset received the undergraduate degrees in computer science and business from the University of Pennsylvania, Philadelphia, PA, USA, and the M.S. and Ph.D. degrees in mechanical engineering from Caltech, Pasadena, CA, USA.

He is a Professor in the Robotics Institute, Carnegie Mellon, Pittsburgh, PA, USA.



Hossein Faraji received the Ph.D. degree in robotics with a focus on dynamics, control, and geometric mechanics from Oregon State University, Corvallis, OR, USA.

He is a Robotics and Control Researcher and Engineer. His research interests include control, motion planning, and calibration of the medical robotics systems.



Nathan Justus received the B.S. degree in aerospace engineering from the University of Oklahoma, Norman, OK, USA, and the M.S. degree in robotics from Oregon State University, Corvallis, OR, USA, where he is working toward the Ph.D. degree in robotics.

He was previously an ISS flight controller with NASA, Washington, DC, USA.



Ruijie Fu received the bachelor's of science degree in mathematics, in 2020, from Carnegie Mellon University, Pittsburgh PA, USA, where she is currently working toward the M.S. degree in robotics.

Her research interests include differential geometry and robot localization.



Suresh Ramasamy received the Ph.D. degree in robotics from Oregon State University, Corvallis, OR, USA, the MS degree in mechanical engineering from Carnegie Mellon University, Pittsburgh PA, USA, and the B.Tech degree in mechanical engineering from the Indian Institute of Technology, Gandhinagar, Gujarat, India.

He is a Senior Data Scientist with CVS Health, Woonsocket, RI, USA.