

Kinematic Cartography for Locomotion at Low Reynolds Numbers

Ross L. Hatton and Howie Choset
Carnegie Mellon University
{rlhatton, choset}@cmu.edu

Abstract—Kinematic motion planning often requires a notion of “distance” between configurations. Euclidean distances on a parameter space are easy to compute, but can drastically distort the effort required to change configuration. Here, we present a framework for characterizing this distortion, based on principles adopted from the cartographic community, and a method for transforming configuration coordinates to better represent actuation costs. As a demonstration of this approach, we derive a true configuration distance metric for an important class of locomoting systems: low Reynolds number swimmers. Applying our cartographic coordinate transformation to these systems both provides intuition for previous numerical results, and allows direct geometric comparison between systems with heterogeneous morphology.

I. INTRODUCTION

In motion planning and locomotion analysis, it is often useful to consider the “distance” between two configurations or the “length” of a given trajectory through the configuration space. These quantities are often defined by a simple Euclidean metric on the system parameters, but several limitations make this approach unsatisfactory for many applications. First, it inherently requires the determination of scaling factors to relate heterogeneous configuration parameters, such as translation and rotation. Second, it includes an often-false underlying assumption that the costs for changing configuration are both uniform and orthogonal – *i.e.*, that they depend neither on the current configuration nor on any coordination between changes in different parameters.

These problems bear a strong resemblance to map projection issues considered in theoretical cartography, a domain of differential geometry that has evolved largely independently from those of geometric mechanics and robotics. Drawing on established practices from the cartographic community [1], [2], we have developed a means for representing how given parameterizations distort effort metrics on the configuration space and then reparameterizing the shape spaces to optimally flatten these metrics. Such reparameterization ensures that the shape space is faithfully represented for plotting or sampling purposes, in much the same way that certain maps are more accurate than others at representing the surface of the globe.

As a demonstration of this cartographic approach, we apply it to an important class of systems: low Reynolds number swimmers. This class incorporates most micro-organisms, and has been the target of increasing interest in the robotics [3], biological [4], and applied mechanics [5]–[7] communities. For these systems, we first introduce a power-based con-




Fig. 1: Model coordinates

figuration distance metric that directly corresponds to the energy dissipated while changing shape. Working with this metric allows for true comparison of configuration distances and trajectory lengths, abstracting away the need to consider pacing along trajectories.

We finish the paper by taking the geometric study of locomotion beyond simple articulated systems and apply it to systems with continuous curvature, representing either truly flexible systems or highly-articulated systems coordinating their joints into fundamental deformation modes. By applying the dissipation metric and cartographic reparameterization, we directly compare the locomotive efficiencies of articulated and continuous swimmers both in general and for specific optimized strokes.

II. GEOMETRIC MECHANICS FOR LOCOMOTION

The analysis in this paper draws on a geometric mechanics locomotion framework that has rigorous theoretical roots in differential geometry [8], but in our work, we use the vector calculus terminology introduced in [9], [10] for broader accessibility.

A. The Reconstruction Equation and the Local Connection

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 M , such that the position $g \in G$ locates the system in the world, and the shape $r \in M$ gives the relative arrangement of the particles that compose it.¹ For example, the position of the three-link system in Fig. 1(a) is the location and orientation

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

of the middle link, $g = (x, y, \theta) \in SE(2)$,² and its shape is parameterized by the two joint angles, $r = (\alpha_1, \alpha_2)$.

With this separation, locomotion is readily seen as the means by which changes in shape (such as strokes, gaits, or wingbeats) affect the position. The geometric mechanics communities in robotics and physics [8]–[14] have studied this process with the development of the *reconstruction equation* and the *local connection*, tools for relating the body velocity of the system, ξ , *i.e.* its longitudinal, lateral, and rotational velocity as depicted in Fig. 1(b), to its shape velocity, \dot{r} , and accumulated momentum.

For systems that are sufficiently constrained or unconstrained, the generalized momentum drops out and the system behavior is dictated by the *kinematic reconstruction equation*,

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

in which the local connection \mathbf{A} acts as a kind of Jacobian [10], mapping from velocities in the shape space to the corresponding body velocity. This kinematic condition has been demonstrated for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [14], [15], non-holonomic constraints such as passive wheels [8], [13], [14], [16], and fluid interactions at the extremes of low [17]–[19] and high [18]–[21] Reynolds numbers. Experimental results also suggest that sandswimming [22] is largely kinematic, though locomotion in this regime has yet to be incorporated into the body of geometric mechanics literature.

B. Low Reynolds Number Swimming

In this paper, we are particularly interested in locomotion at low Reynolds number, *i.e.*, where drag forces dominate inertial forces. This situation most commonly arises for systems swimming at very small scales, but also describes some terrestrial locomotion systems with viscous friction. To derive the local connections for these systems, we draw on the resistive force model we described in [19], based on [6], [7], [18] and summarized below.

At their heart, geometric models for low Reynolds number locomotion are built on two key assumptions. First, that the drag forces are linearly related to the system’s body and shape velocities. Second, that the large drag forces settle the system to its terminal velocity on a much faster time scale than that of any inertial accelerations, so the net forces on the system can be taken as zero for all time. Together, these conditions impose a *Pfaffian constraint* on the system’s generalized velocity,

$$\begin{bmatrix} F_x^b \\ F_y^b \\ F_\theta^b \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \omega(r) \begin{bmatrix} \xi \\ \dot{r} \end{bmatrix}, \quad (2)$$

in which the matrix ω that maps the velocities to the net forces on the body frame depends only on the shape r . By separating

ω into two sub-blocks, $\omega = [\omega_1^{3 \times 3}, \omega_2^{3 \times 2}]$, it is straightforward to rearrange (2) into

$$\xi = -\omega_1^{-1}\omega_2\dot{r}, \quad (3)$$

revealing the local connection as $\mathbf{A} = \omega_1^{-1}\omega_2$.

The ω matrix is generated by calculating the local drag forces on each portion of the body, then transforming these forces into force-moment pairs acting on the body frame and extracting the coefficients corresponding to each component of the generalized velocity. In general, these drag forces are found by solving for Stokes flow around the swimmer; for slender bodies, it is convenient to use Cox theory [23], which first finds the drag on the system as if the particles in the body were independently moving through the fluid, and then iteratively incorporates the wake-coupling interactions between the particles.

In practice, these wake-coupling effects only play a major role when the slender body folds in tightly on itself. As such tightly folded shapes turn out to be inherently inefficient for locomotion, the higher-order terms can be neglected, leaving only the first-order resistive-force terms. The relative drag coefficients in the directions tangential and normal to the swimmer’s body are a function of the slenderness ratio, with the normal coefficient approaching a maximum value of twice the tangential coefficient as the swimmer approaches zero thickness. An example of integrating these forces to find the local connection for a three-link low Reynolds number swimmer is given in [19].

C. Integrating Motion

The reconstruction equation (1) is a differential equation relating the evolution of the system’s position to its trajectory through shape space. This equation can be interpreted as a process model for simulating locomotion, but its real power lies in the way the structure of the local connection captures the locomotive capabilities of the system. In particular, the net displacement for a gait (a closed trajectory in the shape space) can be closely approximated by the area integral of the *curvature form* (Lie bracket) [20] of the local connection over the region of the shape space enclosed by the gait. Figure 2 shows two such gaits, overlaid on the x component of the curvature form corresponding to the low Reynolds number dynamics discussed above. Historically, the error dynamics of this approximation meant that it was regarded as only useful for small-amplitude gaits or certain special cases [20]. In our recent work [19], [24], however, we have demonstrated that optimizing the choice of coordinates can significantly improve the error performance, and successfully applied the area rules to gaits with joint motions in excess of $\pm\pi$ radians.

D. Optimal Strokes

The curvature functions provide some key insights into optimal stroke patterns for a low Reynolds number swimmer found in [7]. Figure 2 shows the maximum-displacement and -efficiency strokes found in that paper plotted against the x

² $SE(2)$ is the set of all translations and rotations in the plane.




Fig. 2: The efficiency-optimal stroke for the three-link swimmer found in [7], overlaid on a contour plot of the x row of the curvature of the local connection. Over each gait, the net displacement per cycle is the integral of the curvature over the enclosed region.

component of the swimmer’s constraint curvature. Immediately, we see that the maximum-displacement stroke follows the zero-contour of the function, thus maximizing displacement integral. The optimal-efficiency stroke cuts across the ends of the negative region (the inclusion of which would increase the displacement per cycle), exchanging these low-yield regions for a shorter stroke length; it is more efficient to execute multiple high-yield cycles than to significantly increase the cost of each cycle for little gain. Similarly, this stroke cuts very slightly into the positive regions of the x function; the slight positive inclusion not being significant enough to justify the extra path-length in the shape space required to completely avoid the positive region.

III. POWER-NORMALIZED COORDINATES

In the discussion of the optimal stroke in §II-D, the cost of the stroke is implicitly taken as the path-length of the stroke through the joint-angle space of the system. While this metric appears reasonable for a symmetric system with identical joints, closer consideration immediately raises serious questions with regard to its suitability: Does the cost to change shape depend on the current shape? Are the joint costs truly independent, or does the motion of one joint help or impede the motion of the other? How can the costs to change shape of two systems with different geometries be compared?

These questions bear strong resemblance to the fundamental question in cartography: Does a given map projection accurately represent the world? The first two questions correspond respectively to how well the projection preserves relative scale across the map and how well it minimizes distortion and preserves angles. The cartographical analog of the last question is less commonly encountered in practice, but would appear in a comparison of maps of dissimilar objects, such as planets of different size or irregularly-shaped asteroids.

Cartographers have developed powerful tools to address these issues and identify accurate map projections. Applying them to our locomoting systems provides shape-space reparameterizations in which path-length closely corresponds to energy usage. In addition to improving the usefulness of the curvature functions in identifying and rationalizing optimal strokes, this normalization directly enables broad efficiency

characterizations between systems with very different morphologies.

A. Distance metrics

The first step to adapting cartographical approaches to our kinematics problems is a basic understanding of distance metrics in mapping. The most commonly encountered distance metric is the familiar Euclidean distance D between two points in \mathbb{R}^2 (the uniform plane),

$$D_{\mathbb{R}} = \sqrt{\Delta x^2 + \Delta y^2}. \quad (4)$$

This metric is a special case, in that the vector-space nature of the plane allows it to be explicitly defined for points a finite distance apart. A more general way to define Euclidean distance is to relate differential changes in x and y to the resulting differential arclength ds traveled by the system,

$$ds_{\mathbb{R}} = \sqrt{dx^2 + dy^2}, \quad (5)$$

and then take distance as the arclength of the shortest path between two points. Such shortest paths are termed *geodesics*, representing the “straightest” lines through the space (and, for this Euclidean example, are straight in the commonly understood sense).

In differential geometry, ds is referred to as the *line element* or “unit infinitesimal length element” for the Euclidean plane. Its square is the *first fundamental form* of the space, and can be represented quadratically as

$$ds_{\mathbb{R}}^2 = dx^2 + dy^2 = \underbrace{\begin{bmatrix} dx & dy \end{bmatrix}}_{\mathcal{M}_{\mathbb{R}}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} dx \\ dy \end{bmatrix}, \quad (6)$$

where $\mathcal{M}_{\mathbb{R}}$ is the *metric tensor* corresponding to the Euclidean metric. Its status as an identity matrix reflects the fact that dx and dy contribute in equal and unit amounts to ds .

Measuring distances on a sphere is less straightforward. The meridians converge near the poles, causing the influence of longitude on the position to drop away as the cosine of the latitude. Although it is possible to define a Euclidean-distance-style metric on the sphere (the great-circle distance), for the purposes of studying map distortion it is more useful to consider the local distance metric. For the unit sphere \mathbb{S}^2 , parameterized by longitude λ and latitude ϕ , the line element is

$$ds_{\mathbb{S}} = \sqrt{(\cos(\phi)d\lambda)^2 + d\phi^2}, \quad (7)$$

and the first fundamental form and metric tensor are thus

$$ds_{\mathbb{S}}^2 = \begin{bmatrix} d\lambda & d\phi \end{bmatrix} \underbrace{\begin{bmatrix} \cos^2(\phi) & 0 \\ 0 & 1 \end{bmatrix}}_{\mathcal{M}_{\mathbb{S}}} \begin{bmatrix} d\lambda \\ d\phi \end{bmatrix}. \quad (8)$$

Great circles are the geodesics defined by this metric: the section of a great circle connecting two points on the sphere minimizes $\int ds_{\mathbb{S}}$ over all possible paths between them, and thus defines the distance between them.

B. Distance metrics in the shape space

Returning to mechanical systems, the most straightforward way to measure the distance between shapes or the “length” of a shape change trajectory is via the Euclidean metric on the space of joint angles,

$$ds_\alpha^2 = d\alpha_1^2 + d\alpha_2^2 = \underbrace{\begin{bmatrix} d\alpha_1 & d\alpha_2 \end{bmatrix}}_{\mathcal{M}_\alpha} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} d\alpha_1 \\ d\alpha_2 \end{bmatrix}. \quad (9)$$

As previously noted, this metric only considers joint motion, and not the effort required to produce that motion. Happily, for many systems, we *can* define distance metrics on the shape space that directly correspond to such effort – for example, the inertia tensor of a mechanical system with Lagrangian dynamics induces a metric whose geodesics correspond to minimum energy paths between configurations [25].

At low Reynolds numbers, the dissipation forces from fluid drag are much larger than any inertial forces on the system, and the inertia tensor is thus not a physically meaningful configuration metric for systems in viscous regimes. The rate-dependent nature of these dissipative forces has led previous investigations of low Reynolds number swimming (*e.g.*, [6], [7]) to calculate the cost of motion along a given shape trajectory as a pacing-dependent time integral, with a secondary calculation to determine the optimal pacing. We have found, however, a means of turning the dissipation forces into a configuration metric in which arclength is a *pacing-independent* measure of the effort required by a shape trajectory, and whose geodesics correspond to minimum-effort changes in shape.

To find this metric, we start by noting that the linear relationship between force and velocity that underlies (2) also generates a linear relationship between the torques \mathcal{T} on each joint and the generalized velocities. Incorporating the relationship between body and shape velocities from (1) allows the elimination of the body velocity terms in this equation, leaving the torque as

$$\mathcal{T} = f(\alpha)\dot{\alpha}, \quad (10)$$

where f is a symmetric 2×2 matrix whose entries are functions of the system shape. The power p dissipated through the joints is the dot product of their velocities and torques, $p = \dot{\alpha} \cdot \mathcal{T}$, or, combined with (10),

$$p = \dot{\alpha}^T f(\alpha) \dot{\alpha}. \quad (11)$$

The quadratic nature of (11) bears a strong resemblance to the metric tensor formulas above. By separating $\dot{\alpha}$ into time and angular differentials and rearranging terms, we can bring (11) explicitly into this form as

$$ds_p^2 = p dt^2 = \underbrace{\begin{bmatrix} d\alpha_1 & d\alpha_2 \end{bmatrix}}_{\mathcal{M}_p} f(\alpha) \begin{bmatrix} d\alpha_1 \\ d\alpha_2 \end{bmatrix}, \quad (12)$$

where f now takes on the additional role of a metric tensor, encoding the effort required to move through the space. The

line element for this space,

$$ds_p = \sqrt{p} dt, \quad (13)$$

initially appears somewhat odd, containing both time and the square root of power in a length definition. It does, however, have an intuitive physical interpretation: reaching points that are farther away either takes longer or requires moving at increased speed, and power is proportional to velocity squared.

A key benefit to considering shape changes in the power-dissipation coordinates is that it abstracts out much of the need to consider pacing within a stroke. Becker *et al.* [6] showed that for any given cyclic frequency, the optimal pacing is the one in which the dissipated power is held constant. This optimal pacing directly corresponds to moving through the shape space with constant speed ds_p/dt . Thus, by measuring the input costs of strokes by their $s_p = \sqrt{p}(\Delta t)$ length, we can make fundamental, pacing-free comparisons of their input costs.

C. Tissot’s indicatrix

Equirectangular map projections, such as that in Fig. 3, take the latitude and longitude as evenly spaced points on a grid. In doing so, they implicitly apply the Euclidean metric in (6) to points on the sphere, in place of the natural spherical metric in (8). This exchange of metrics severely impacts our perception of the world – while we can of course *compute* distances and path-lengths, straight lines on the map are not the shortest paths on the sphere, and different regions of the map are stretched out of proportion with each other. Similarly, the distance metric in (12) corresponds to a (generally) non-flat manifold.³ Representing the shape space of the three-link swimmer as an (α_1, α_2) grid, as in the plots in Fig. 2, thus distorts both the shapes of the functions and the relationship between effort and pathlength.

Tissot’s indicatrix [1] illustrates the extent of these distortions by showing how they transform (infinitesimal) circles into ellipses at different points in the parameter space. In the equirectangular projection in Fig. 3, there is no distortion at the equator and indicatrices there are perfect circles. The distortion increases with the latitude, stretching the indicatrices (and the map features) longitudinally. Note that this stretch is *inversely* proportional to the metric – as the top left entry of \mathcal{M}_S becomes smaller, it “costs less distance” to change the longitude by a given amount.

We calculate the indicatrices from the singular value decomposition of the metric tensor. The metric tensor is symmetric, so this operation takes the form $\text{svd}(\mathcal{M}) = U\Sigma U^T$, where the columns of U are unit vectors encoding the *principle stretch directions* [2], *i.e.*, the directions in which the distance metric is longest and shortest. The values in the diagonal matrix Σ are the magnitudes of ds^2 along the corresponding columns of U , and thus are the inverse squares of the *principle stretches* τ , the magnitudes of deformation along these axes. From these

³Roughly, a space that is itself curved through higher dimensions.

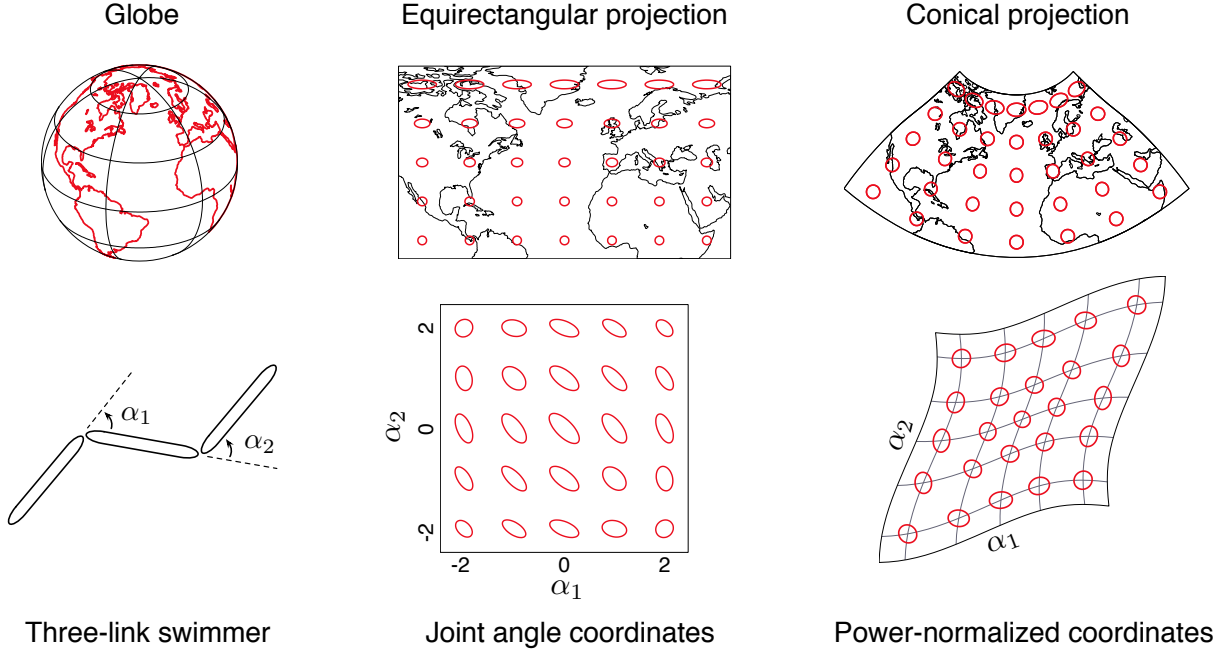


Fig. 3: Tissot’s indicatrices for the world map and three-link swimmer’s power dissipation metric, shown in both the original and relaxed coordinates.

properties, we can define a *Tissot transformation matrix* T of a metric,

$$T = U\Sigma^{-\frac{1}{2}}U^T, \quad (14)$$

that encodes the local deformation produced by the projection. We should note here that although the cartographic literature uses a similar approach to finding the principal stretches, we have not been able to find examples where T is explicitly constructed, with the authors in that field content to note that U and $\Sigma^{-\frac{1}{2}}$ respectively give the directions and magnitudes of the semi-axes of the indicatrices.

Applying Tissot’s indicatrix to the three-link swimmer, as in the middle-bottom plot in Fig. 3 shows how distances calculated according to the power metric \mathcal{M}_p are distorted when the system is represented in joint angle coordinates: First, the major axes of the indicatrices are all in the $(+\alpha_1, -\alpha_2)$ direction, indicating that this representation generally stretches features along this axis (or, equivalently, compresses them in the orthogonal direction). Second, the indicatrices are generally larger near the center of the space, where one or both of the joint angles are zero. Physically, these observations correspond to the system meeting less resistance from the fluid when moving the joints together, rather than oppositely; even actions or those starting from large joint angles tend to push the links laterally (in the high-drag direction); in odd motions, the drag torques induced in response to a given joint tend to “help” the motion of the other joints.

D. Optimal Projections

Altering a map projection to de-identify the latitude and longitude from the xy coordinates of the paper can significantly reduce the distortion. Conical projections, such as the one at

the top right of Fig. 3, preserve distances and orthogonality over much of their domain, as reflected in the uniform size and near-circularity of the Tissot indicatrices. Distortion appears only very close to the singularity at the north pole, and is much less severe than in the equirectangular projection.

Several metrics are available for evaluating projections, based on how well they preserve linear distance, orthogonality, and area. Of these, there are strong cartographical arguments [2] for using linear distance as a primary criterion for map quality; such a metric also fits well with our goal of finding a projection of the shape space that best preserves the distance metric in (12). There are two general categories of linear-distance criteria [2]: *Tissot criteria*, based on the ratio between largest and smallest indicatrix axes in the domain, and *Airy criteria*, that consider the average squared stretch over the domain.

Cartographers optimizing maps for these criteria typically make use of the (oblate) spherical nature of the earth. The mathematical “purity” of this shape makes it relatively straightforward to project from the globe onto a cylindrical or conical surface, and then to optimize the placement of this surface to minimize the distortion. The natural form of the shape space, defined by the power metric \mathcal{M}_p , is considerably more complicated and does not readily admit such geometrical approaches. Instead, we have developed a numerical approach for finding the change of coordinates that best inverts the stretches illustrated by Tissot’s indicatrix.

In this approach, we treat the shape space as an elastic sheet with pre-existing internal strains (the Tissot stretches) and allow it to relax into a minimum-energy configuration. This relaxation corresponds to optimizing an Airy criterion, in that the strain energy in the sheet corresponds to the squared

TABLE I: Distortion metrics

System	Tissot Criterion crit_T	Jordan Criterion crit_J	equiv. ratio
Map, equirectangular	5.3367	0.1355	1.6943
Map, conical	1.7981	0.0031	1.0930
Swimmer, joint angles	3.2863	0.0667	1.4673
Swimmer, power-normed	1.2109	0.0094	1.1643

stretch. As the sheet relaxes, points specified in the starting coordinates (*e.g.*, latitude and longitude or the joint angles) are pulled into new positions in which the inter-point distances are as consistent with the preferred distance metric (*e.g.*, distance on a sphere, power dissipation) as is possible while remaining in the plane. Taking the xy coordinates of this plane as our new parameterization then gives a basis that better matches the metric, which we can convert into our original coordinates as needed by inverting the relaxation transform.

Applying this procedure (the details of our implementation are given below) to a map initially plotted as an equirectangular projection pinches in the pole and curves the lines of constant latitude in the same manner as the conical projection, and the “conical projection” in the top right image of Fig. 3 is actually the product of our relaxation algorithm. For the swimmer, coordinate relaxation has the effect of “scissoring” the primary α axes together, extending the first and third quadrants while minimally affecting the second and fourth quadrants, as shown at the bottom right of Fig. 3.

Table I quantifies the efficacy of these coordinate changes by means of both the Tissot criterion for maximum stretch ratio, $\text{crit}_T = \tau_{1,\text{max}}/\tau_{2,\text{min}}$ and the Jordan form [2] of the Airy criterion for average stretch, the details of which are in the appendix. Included with those criteria are the constant indicatrix aspect ratio that would produce the same Jordan value. For the swimmer, relaxing the parameterization has reduced the maximum stretch in the figure from 229% to 21% and the average stretch at any location from 47% to 16%.

We should note that some previous efforts, in particular [17], have made preliminary use of a dissipation-based metric on the shape space for low Reynolds number swimming. In that work, however, the change from joint angles to power-based distance measures was only made locally, by (in our terminology) multiplying the constraint curvature by $\det(T^{-1})$ at each pair of joint angles. This approach works for differential oscillations, but the variations in T across the space (which we balance via our elastic relaxation process) prevent it from being a viable means of representing large-scale locomotive efficiency.

E. Numerical approach

We implement the coordinate relaxation process on an eight-connected network of springs evenly distributed over the starting parameter space. This network naturally decomposes into cells containing four nodes and six springs, as illustrated in Fig. 4(a). Within each cell, we find the natural length of the springs by approximating the deformation as constant




Fig. 4: Components of the numerical algorithm.

over the cell and transforming the nodes by the inverse of the Tissot transformation matrix: as shown in Fig. 4(b), the nodes have been brought onto the local unit circle from their natural positions (according to the preferred metric) by the same transformation T that maps the unit circle to Tissot’s indicatrix. Inverting this transform reverts them to their natural positions, where the inter-node distances give the resting spring lengths (*i.e.*, a linearized approximation to the metric-defined distances between the nodes). For springs that are shared between two cells, we average the lengths found.

Once the spring lengths have been found, a variety of methods are available for minimizing the residual energy. Our results here are based on assigning equal linear stiffnesses to each spring and linear damping at the nodes, then integrating forward in time until the system reaches equilibrium. Refinements to this approach that we are presently investigating include finite element methods to directly locate equilibria, nonlinear springs, and emphasizing targeted regions of the parameter space via variations in stiffness.

F. Optimal stroke in new coordinates

Figure 5 shows the optimal strokes from [7] plotted in the power normalized shape coordinates r_p , along with the x component of the curvature function. Immediately, we can see that the cost savings in the maximum-efficiency stroke (as compared to the maximum-displacement stroke) are much greater than they appeared in the original coordinates: the difference is along the even axis, and so accounts for a proportionally larger share of the arclength in the new coordinates. Interestingly, the double-well feature previously noted in the x curvature function is absent here, suggesting that it was an artifact of the joint-angle parameterization.

This representation also provides some intuition behind the constant-power nature of optimal pacing: slowing down any one segment requires a corresponding increase in speed elsewhere in the stroke. As the distance metric is quadratic in power, the increase will always cost more than was gained by slowing down, with little opportunity for arbitrage by speeding up in a low-cost region.

IV. OTHER SHAPE MODES

The geometric principles outlined above do not depend on the three-link nature of the swimmer, and are applicable for any shape parameterization (including flexible systems

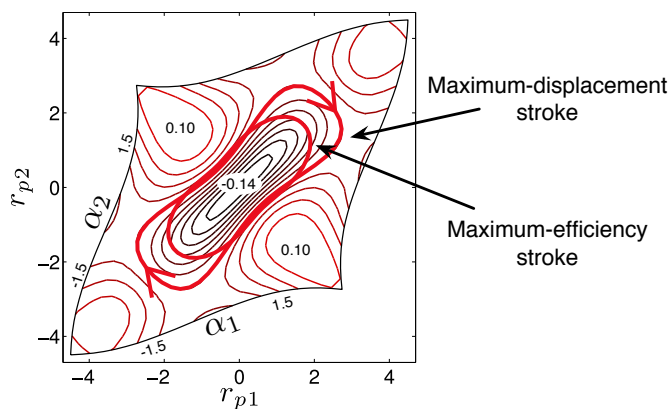


Fig. 5: Optimal strokes and x curvature in power-normalized coordinates.

and highly-articulated systems with coupled joint angles); the prominence of three-link systems in the literature is chiefly a matter of convenience. For systems with disparate shapes, comparisons of locomotive efficiency have largely relied on first finding the optimal strokes for each system – the differences in power costs creating difficulty in making broader observations. Our power-normalized coordinates open a new avenue for addressing this limitation: normalizing the coordinates to matching scales encodes the efficiency comparisons directly in the topography of the constraint curvature functions.

While a full investigation of this principle is beyond the scope of this paper, a simple application provides some key insight into the advantages swimming mechanisms with smooth curvature have over those with rigid links. For the purposes of comparison, we will consider a system whose shape is determined by two single-period sinusoidal curvature modes along the length of the backbone,⁴ as shown at the right of Fig. 6. The shape parameters a for this system correspond to the amplitudes of the curvature bases κ , just as the joint angles of the three-link swimmer multiply the impulse-function curvature bases that define an articulated system. The overall shape produced by the net sinusoidal curvature is a *serpenoid curve* [26]; this curve is rounder than a sinusoid and has been demonstrated to closely model the motions of snakes and other undulatory locomoting animals [26]. A *serpenoid wave* of constant amplitude and phase velocity is the most common dynamic application of such curves, and corresponds to tracing out a circle in the (a_1, a_2) space with constant speed.

Power-normalizing the serpenoid amplitude coefficients, as in Fig. 7, has similar results to those seen for the three-link swimmer – the even mode (κ_1) costs more to actuate than the odd mode (κ_2), and the sign-definite region at the center of the x constraint curvature function is longer in the even direction than the odd direction. Comparing the magnitudes of the curvature functions, however, reveals the advantage held by the serpenoid swimmer. First, the x curvature for the continuous modes has a peak magnitude more than twice that of the rigid-link swimmer, indicating that the serpenoid swimmer is considerably more efficient when locomoting via

⁴Note that this (physical) curvature of the backbone is distinct from the (abstract) curvature of the constraint functions.




Fig. 6: Curvature modes and their corresponding shapes.




Fig. 7: Power-normalized projection of the serpenoid swimmer's x constraint curvature and a constant-amplitude serpenoid wave. Tissot distortion in this plot is 1.21, Jordan-equivalent distortion is 1.07.

differential oscillations. Second, the optimal strokes for this serpenoid curve are significantly more compact than the rigid-link stroke, providing an improved perimeter-to-area ratio.

Stroke lengths and net x displacements for the three-link and serpenoid swimmers are compared in Table II. These numbers for unit drag on body lengths of six units, corresponding to the unit link half-length in [7], [19], and scale with length as indicated in the table. The stroke presented for the serpenoid system is the serpenoid wave that just crosses into the negative regions of the curvature function; this stroke is optimal over the serpenoid waves, though small performance increases could likely be gained by using a higher-order stroke model as in [7]. Even without such optimizations, however, the serpenoid swimmer significantly outperforms the three-link system, generating over three times the displacement per unit of effort. Physically, this corresponds to the rigid system taking three times as long to cross a given distance at equal power (and thus three times as much energy), or *nine* times the power and energy to travel the distance in the same time as the serpenoid swimmer.

A second interesting observation regarding the serpenoid stroke in Fig. 7 is that the compression of the odd axis suggests the use of non-uniform pacing in the serpenoid wave. Following the ellipse with constant speed corresponds to accelerating the wave when the odd mode is changing, taking advantage of the lower fluid resistance to this motion. We believe that such pacing questions have not yet been investigated in the serpenoid literature, opening another avenue of future work.

TABLE II: Stroke efficiencies

Quantity	Three-link	Serpentoid	Length scaling
Stroke length s_p	24.26	9.33	L^3
Net x displacement	0.61	0.85	L
Displacement per s_p	0.03	0.09	L^{-2}

V. CONCLUSIONS AND FUTURE WORK

To the best of our knowledge, this paper is the first to apply cartographic distortion analysis to kinematics and locomotion. By unifying these two branches of differential geometry, we have found coordinate representations that usefully encode notions of distance and effort. These coordinates increase the tractability of several motion planning questions and enable simple and direct comparison of distinctly heterogeneous locomoting systems. In developing these coordinate representations, we introduced a new metric for low Reynolds number actions that allows pacing-free comparisons of the effort required to move along different trajectories.

Looking forward, we see many opportunities to extend the paradigm introduced here. This work will include an expansion of the stroke comparisons in §IV (considering both richer shape modes [27] and trajectories [7]), investigations into distance metrics for other physical regimes or that incorporate internal dynamics [28], and applications to other areas, such as even sampling of the shape space or proper computation of system manipulability metrics. Alongside this work, we will continue to develop its theoretical underpinnings, especially the effects of using different metrics in the elastic relaxation process, implementations in higher dimensions, and a more fully differential-geometric presentation of the principles.

We thank Anette Hosoi, Lisa Burton, Matthew Tesch, and Glenn Wagner for many insightful discussions, and the NSF for support via Grant 1000389.

APPENDIX: JORDAN CRITERION

The Jordan criterion [2] measures the average linear distortion of a projection in all directions, and is calculated as

$$\text{crit}_J = \frac{1}{A} \iint_{\Omega} \left(\frac{1}{2\pi} \int_0^{2\pi} (R(\Theta) - 1)^2 d\Theta \right), \quad (15)$$

where Ω is the region of area A under consideration, and

$$R(\Theta) = \tau_1 \tau_2 / \left(\text{mean}_{\Omega}(\tau) \sqrt{\tau_1^2 \sin^2 \Theta + \tau_2^2 \cos^2 \Theta} \right) \quad (16)$$

is the radius of the Tissot indicatrix at an angle Θ from the major axis, normalized by the mean stretch over the domain.

REFERENCES

- [1] M. Tissot, *Memoire sur la Representation des Surfaces et les Projections des Cartes Geographiques*. Paris, France: Gautier-Villars, 1881.
- [2] A. Niermann, "A comparison of various advantageous cartographical representations," Master's thesis, Universität Stuttgart, Stuttgart, Germany, June 1983.
- [3] K. Peyer, L. Zhang, B. Kratochvil, and B. Nelson, "Non-ideal swimming of artificial bacterial flagella near a surface," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, USA, May 2010, pp. 96–101.
- [4] G. Taylor, "Analysis of the Swimming of Microscopic Organisms," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 209, no. 29, pp. 447–461, 1951.
- [5] E. M. Purcell, "Life at low reynolds numbers," *American Journal of Physics*, vol. 45, no. 1, pp. 3–11, January 1977.
- [6] L. Becker, S. A. Koehler, and H. A. Stone, "On Self-propulsion of Micro-machines at Low Reynolds Number: Purcell's three-link swimmer," *Journal of Fluid Mechanics*, vol. 490, pp. 15–35, 2003.
- [7] D. Tam and A. E. Hosoi, "Optimal stroke patterns for purcell's three-link swimmer," *Phys. Review Letters*, vol. 98, no. 6, p. 068105, 2007.
- [8] A. M. Bloch *et al.*, *Nonholonomic Mechanics and Control*. Springer, 2003.
- [9] R. L. Hatton and H. Choset, "Connection vector fields for underactuated systems," in *Proceedings of the IEEE BioRobotics Conference*, October 2008, pp. 451–456.
- [10] —, "Optimizing coordinate choice for locomoting systems," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, USA, May 2010.
- [11] A. Shapere and F. Wilczek, "Geometry of Self-Propulsion at Low Reynolds Number," *Journal of Fluid Mechanics*, vol. 198, pp. 557–585, 1989.
- [12] S. Kelly and R. M. Murray, "Geometric Phases and Robotic Locomotion," *J. Robotic Systems*, vol. 12, no. 6, pp. 417–431, Jan 1995.
- [13] J. Ostrowski and J. Burdick, "The Mechanics and Control of Undulatory Locomotion," *International Journal of Robotics Research*, vol. 17, no. 7, pp. 683 – 701, July 1998.
- [14] E. A. Shammass, H. Choset, and A. A. Rizzi, "Geometric motion planning analysis for two classes of underactuated mechanical systems," *Int. J. of Robotics Research*, vol. 26, no. 10, pp. 1043–1073, 2007.
- [15] G. Walsh and S. Sastry, "On reorienting linked rigid bodies using internal motions," *Robotics and Automation, IEEE Transactions on*, vol. 11, no. 1, pp. 139–146, January 1995.
- [16] R. Murray and S. Sastry, "Nonholonomic Motion Planning: Steering Using Sinusoids," *IEEE Transactions on Automatic Control*, vol. 38, no. 5, pp. 700–716, Jan 1993.
- [17] J. Avron and O. Raz, "A geometric theory of swimming: Purcell's swimmer and its symmetrized cousin," *New Journal of Physics*, vol. 9, no. 437, 2008.
- [18] S. D. Kelly, "The mechanics and control of driftless swimming," in press.
- [19] R. L. Hatton and H. Choset, "Connection vector fields and optimized coordinates for swimming systems at low and high reynolds numbers," in *Proceedings of the ASME Dynamic Systems and Controls Conference (DSCC)*, Cambridge, Massachusetts, USA, Sep 2010.
- [20] J. B. Melli, C. W. Rowley, and D. S. Rufat, "Motion Planning for an Articulated Body in a Perfect Planar Fluid," *SIAM Journal of Applied Dynamical Systems*, vol. 5, no. 4, pp. 650–669, November 2006.
- [21] E. Kanso, "Swimming Due to Transverse Shape Deformations," *Journal of Fluid Mechanics*, vol. 631, pp. 127–148, 2009.
- [22] R. D. Maladen, Y. Ding, C. Li, and D. I. Goldman, "Undulatory swimming in sand: Subsurface locomotion of the sandfish lizard," *Science*, vol. 325, no. 5938, pp. 314–318, July 2009.
- [23] R. G. Cox, "The motion of long slender bodies in a viscous fluid part 1. general theory," *J. Fluid Mechanics*, vol. 44, no. 4, pp. 791–810, 1970.
- [24] R. L. Hatton and H. Choset, "Geometric motion planning: The local connection, stokes's theorem, and the importance of coordinate choice," *International Journal of Robotics Research*, to appear June 2011.
- [25] S. Sternberg, *Lectures on Differential Geometry*. Chelsea Publishing Company, 1982.
- [26] S. Hirose, *Biologically Inspired Robots (Snake-like Locomotor and Manipulator)*. Oxford University Press, 1993.
- [27] G. Chirikjian and J. Burdick, "A Modal Approach to Hyper-redundant Manipulator Kinematics," in *IEEE Transactions on Robotics and Automation*, vol. 10, 1994, pp. 343–354.
- [28] Z. Guo and L. Mahadevan, "Limbless undulatory propulsion on land," *Proceedings of the National Academy of Sciences*, vol. 105, no. 9, pp. 3179–3184, 2007.