

Geometric Maneuverability

with Applications to Low Reynolds Number Swimming

Ross L. Hatton Lisa J. Burton A. E. Hosoi Howie Choset
rlhatton@cmu.edu lisab@mit.edu peko@mit.edu choset@cmu.edu

Abstract—A mobile system’s *maneuverability* describes the scale and span of the velocities with which it can move. In this paper, we present a new geometric framework for describing the maneuverability of kinematic locomoting systems, inspired by the *manipulability* analysis of robotic arms. This framework describes both the *local maneuverability* in the neighborhood of each shape the system can assume and the *cyclic maneuverability* achieved by executing gaits from a library. Additionally, the gait-level analysis includes tools that direct the search for gaits whose inclusion into the library will usefully improve the maneuverability. Throughout, we provide examples based on a swimming system operating at low Reynolds number.

I. INTRODUCTION

Understanding the maneuverability of mobile systems – the speed and direction with which they move and their ability to change that velocity – is fundamental to their effective design and control. For systems that can apply vector thrust, maneuverability is often well-characterized in terms of top speed and acceleration, but such metrics are less directly applicable to locomoting systems that move by changing shape to push against the environment. The motion capabilities of these systems vary significantly with their shapes, and thus are constantly changing as the system propels itself. Additionally, joint limits restrict the extent of any given action by the system; its motion capabilities over longer distances therefore depend not only on their instantaneous values, but also on the existence of gaits or strokes that efficiently extract net displacement over each cycle.

Recently, we have observed strong similarities between the maneuverability problem for kinematic locomoting systems and the *manipulability* framework widely adopted for the analysis of robotic arms. In this paper, we modify that framework to construct locomotive maneuverability metrics that are at once intuitive and rigorously founded in differential geometry. A relatively direct application of the manipulability principles provides the local maneuverability metric; by incorporating geometric tools for averaging across gait cycles, we generate an equivalent metric for motion over larger length scales. This metric has the additional benefit of directing the search for new useful gaits, by indicating directions in which the search is most likely to be fruitful.

The maneuverability metrics apply to a wide range of kinematic systems and suggest natural extensions to incorporate those with higher-order dynamics. For the purposes of this paper, we demonstrate the maneuverability metrics by applying them to a swimming system at low Reynolds number (*i.e.*, in which viscosity dominates over inertia). This




Fig. 1: Model coordinates

choice reflects both the clear view the swimmer provides of the principles at hand and its applicability to the analyses of both small-scale robotic systems and micro-organisms.

II. GEOMETRIC MECHANICS OF LOCOMOTION

The motion analysis in this paper builds on a foundation of previous work in geometric mechanics, the key aspects of which we highlight below.

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 bodies that compose it.¹ For example, the position of the three-link system in Fig. 1(a) is the location and orientation of the medial line, $g = (x, y, \theta) \in SE(2)$,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 community [1]–[6] has 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*

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*.

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




Fig. 2: The maximum-efficiency and maximum-displacement strokes for the three-link swimmer found in [16], overlaid on a contour plot of the x row of the curvature of the local connection. At right, the strokes are plotted in a cartographically-normalized shape space in which path lengths are proportional to power costs.

equation,

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

in which the local connection \mathbf{A} acts as a kind of Jacobian [5], 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 [4], [7], non-holonomic constraints such as passive wheels [2]–[4], [8], and fluid interactions at the extremes of low [9]–[11] and high [10]–[13] Reynolds numbers. Experimental results also suggest that sandswimming [14] is largely kinematic, though locomotion in this regime has yet to be incorporated into the body of geometric mechanics literature.

B. 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* [12] of the local connection over the region of the shape space enclosed by the gait. Two such gaits appear in Fig. 2, overlaid on the x component of the curvature form corresponding to the low Reynolds number dynamics discussed below. Historically, the error dynamics of this approximation meant that it was regarded as only useful for small-amplitude gaits or certain special cases [12]. In our recent work [11], [15], 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.

III. LOW REYNOLDS NUMBER SWIMMING

Swimmers at low Reynolds number, which include both miniature robots and micro-organisms, form a class of systems that is both the subject of ongoing research across a

range of fields and characteristic of the locomotion principles we are investigating here. These systems move in a regime where the drag forces from the surrounding fluid are linear with velocity and dominate the inertial forces. Their local connections are found by assuming zero net force and torque on the swimmer (a system with Reynolds number $\ll 1$ is so thoroughly drag-dominated that it instantaneously achieves local equilibrium) and finding the combinations of shape and body velocity that balance the drag forces on the swimmer. We describe this process in [11], building on seminal work by Shapere and Wilczek [17], [18] that was further developed by Kelly [10] and Avron and Raz [9].

Local connections and curvature forms for low Reynolds number swimmers appeared as early as [18], which was also one of the earliest works applying differential geometry (in the guise of gauge theory) to any kind of locomotion. The small- θ requirement, however, meant that much of the interesting subsequent work on larger strokes (which tend to be more efficient) was carried out through direct integration of the equations of motion over specific strokes [16], [19], [20], sometimes combined with symmetry arguments. Now that we have separated the small angle requirement from the stroke magnitude, we can use the area integration rules discussed above to graphically and intuitively explain key results from the prior work that bear on the results in this paper.

Tam and Hosoi [16] parameterized the space of strokes for the three-link swimmer as a Fourier series in each joint angle and numerically optimized the coefficients to maximize displacement per stroke (treated in that work as maximizing speed for a given cyclic frequency) and efficiency (displacement per energy expenditure, which is equivalent to maximizing speed at a given power dissipation). Overlaying these optimal strokes with the x component of the curvature form of \mathbf{A} , as in Fig. 2, reveals their genesis: First, the maximum-displacement stroke traces the zero-contour of the x curvature, maximizing the value of the enclosed area integral.

Second, in many cases, the most important gaits are not those that move the farthest in a given cycle, but those that move most efficiently. If we equate “effort” with joint motion, then it is reasonable to expect that the maximum-efficiency stroke will resemble the maximum-displacement stroke, but with a reduced perimeter; this is borne out in Fig. 2, where we see that the optimally-efficient stroke gives up the low-value ends of the negative region in exchange for a significantly shorter sweep of the joints, and thus a better overall displacement-to-cost ratio.

While intuitive, identifying input cost with arclength in the shape space is not completely straightforward, as the power required to move the joints depends on both their present configuration and the speed at which they are moved. For the time-parameterized strokes in [16], these concerns were addressed using the observation in [20] that as the drag forces are linear with respect to the joint velocities, the power dissipation is quadratic. By corollary, the most efficient means of following a path through the shape space

in time T is by moving with constant power; selecting a common cycle time then allows for a power comparison between strokes.

Unfortunately, the power calculation in [16] requires a time parameterization, detracting both from the geometric intuition that underlies other aspects of our analysis and from the notion that a kinematic system’s trajectories should have fundamental properties independent of time. We resolved this discrepancy in [21] by transforming the power metric from [20] into a *distance metric* on the shape space, for which the differential arclength $ds = \sqrt{p} dt$ captures the product of power exerted and time taken to change shape. Moving with a constant speed with respect to this metric directly corresponds to moving with constant power, and arclengths $s = \int \sqrt{p} dt$ are fundamental cost properties of paths through the shape space, independent of the time parameterization: a path with a larger s will always take more power and and/or time to traverse than a path with a smaller value of s .

Projecting the joint angles onto this new distance metric stretches the shape space, as illustrated at the right of Fig. 2, showing the true costs of the two strokes. The dilation of the $\alpha_1 = \alpha_2$ axis emphasizes the reduction in perimeter between the displacement- and efficiency-optimal strokes – the reduction is along the magnified direction, and thus accounts for a proportionally larger contribution of the total stroke cost.

IV. LOCAL MANEUVERABILITY

Using the differential and integral tools outlined above, we can now address the maneuverability of locomoting systems. For the purposes of this discussion, we distinguish between two broad classes of maneuverability: *local maneuverability*, the responsiveness of the system to small joint motions, and *cyclic maneuverability*, the extent to which it can use gait cycles to move effectively over longer distances. These concepts are related respectively to notions of *controllability* [2] and the efficiency metrics in §III, but have some key distinguishing features: local maneuverability considers not only the span of position velocities the system can produce (controllability), but also their magnitudes; cyclic maneuverability applies these concepts of span and magnitude to gait-level motion, expanding on the previous consideration of only the most efficient gait. In many respects, the closest ancestors to our maneuverability measures are the *manipulability* functions commonly applied to robotic arms.

A. Manipulability and Maneuverability

In the analysis of robotic arms, *manipulability* [22] is commonly used as a measure of how rapidly the end effector moves in response to a unit combination of input joint velocities. Manipulability is most commonly expressed either as an ellipse (or ellipsoid) at each configuration, representing the achievable end-effector velocity in each direction for that shape, or as the scalar area (or volume) of the ellipse, representing the average of these velocities across all directions. The manipulability ellipse of the arm is constructed from its

Jacobian,³ J , at each point in its shape space (which, in the absence of a separate position space for a fixed manipulator, is often referred to as the “configuration space”); specifically, for

$$U\Sigma V^T = \text{svd}(J(r)), \quad (2)$$

the columns of U provide the major and minor axes of the ellipse, *i.e.*, the directions in which the end effector can move most and least quickly, and the singular values (drawn from the diagonal elements of Σ) give the magnitudes along these axes. The scalar manipulability value W is the product of these singular values, and can be directly computed as

$$W(r) = \sqrt{\det(J(r)J(r)^T)}, \quad (3)$$

where squaring J and taking the square root of the determinant accommodates systems with redundant degrees of freedom.

Applying these ideas to our locomoting systems, we observe that the local connection in (1) strongly resembles the Jacobian of a manipulator: it maps velocities in the shape space to their resulting velocities in the position space, which, if we consider this our “task space,” is directly equivalent to the end-effector space of an arm. Accordingly, we now propose the local maneuverability of a kinematic locomoting system to be its manipulability, calculated as

$$U\Sigma V^T = \text{svd}(\mathbf{A}(r)). \quad (4)$$

To implement this proposition, two concerns must be addressed. First, the general manipulability framework assumes that the column-space of J is at least equal to the end effector space, so the ellipsoid defined in (2) has a properly defined volume. For some locomoting systems, such as the three-link swimmer and other planar systems with two actuators, the local connection can have more rows than columns, leading to degenerate maneuverability ellipsoids flattened into two-dimensional maneuverability ellipses embedded in a three-dimensional position space; the consequent redundant singular value produces a scalar maneuverability of $W = 0$, despite the system’s clear ability to move itself in some directions. Second, even if J is of full position-rank, computing W on $SE(n)$ at least requires properly scaling the rotational and translational components within the multiplication, and discards their fundamental differences.

While the manipulability literature contains several means of calculating W in spite of the mixed coordinates [23], [24], we have found that for the purposes of describing maneuverability, a collection of non-volume characteristics of the ellipse defined in (4) present a more useful picture of the system capabilities, as illustrated in Fig. 3. To start, we take the *translational maneuverability* as the projection of the full ellipsoid onto the xy plane (or, equivalently, compute the ellipse corresponding to the xy rows of \mathbf{A}). From this ellipse, we can easily extract both the principal and co-principal translational maneuverabilities σ_{xy} (the major and minor semi-axes of the ellipse) and their product, the

³The differential map from joint velocities to end effector velocities.




Fig. 3: Maneuverability ellipsoid for a system with an $SE(2)$ position space, along with its projections and axis intersections. The projections represent the extrema of the ellipse along the translational and rotational directions, while the intersections are the maximum achievable pure-translation or pure-rotation velocities.

scalar translational maneuverability, W_{xy} . Similarly, we can perform the same calculations on the intersection of the ellipsoid with the xy plane to compute the *pure-translation maneuverability*, with components $\sigma_{\overline{xy}}$. These terms are analogous to maximum-speed metrics for traditional vehicles.

Following the same approach, the rotational and pure-rotational maneuverabilities, σ_θ and $\sigma_{\overline{\theta}}$, are naturally identified with the projection and intersection of the ellipse with the θ axis, and correspond to the maximum turning rates of the system. Points on the surface of the ellipsoid represent maximally-fast combinations of rotation and translation, with turning radii inversely proportional to the slope between the θ axis and the lines connecting them to the origin.

Separating the position space into translational and rotational components also solves the problem of working with Jacobians of less-than full rank, as it now does not matter that some components of the maneuverability go to zero. In general, a two-column local connection on $SE(2)$ will produce one of three situations: (a) rotations are coupled to translations, in which case the ellipsoid flattens to an ellipse that obliquely intersects the xy plane, with $\sigma_{\overline{xy}}^2$ and $\sigma_{\overline{\theta}}$ equal to zero; (b) translations are coupled together with $\xi_x = k\xi_y$, so the ellipse is orthogonal to the xy plane with $\sigma_{\overline{xy}}^2$, and $\sigma_{\overline{\theta}}^2$ zero; or (c) there is no rotation, the ellipse lies in the xy plane, and $\sigma_{\overline{xy}}^1 - \sigma_{\overline{xy}}^1$, $\sigma_{xy}^2 - \sigma_{\overline{xy}}^2$, $\sigma_{\overline{\theta}}$, and σ_θ go to zero. In case (a), the turning radius for the system is the slope of the plane containing the ellipse, while in case (b), the system has a choice of turning radii, and in case (c) has, of course, no ability to turn.

B. Choice of Coordinates

An interesting choice appears when calculating the maneuverability of a given system – which reference point on the system should be used? One natural choice, given our previous work, is the optimized coordinate set that we identified in [5], [15]. Among other interpretations, these coordinates maximally decouple changes in position and shape, and so identify a “location” for the system that is as independent of the current shape as possible. The optimal coordinates are based on the system *constraints* and deliberately abstract away its physical components. Here, however, it is precisely the motion of the components with




Fig. 4: Maneuverability ellipsoids (flattened into two-dimensional ellipses oriented in three-dimensional space) for the center of mass and mean orientation line of the three-link swimmer in Fig. 1(a), operating at low Reynolds number. At left, the maneuverability ellipsoids are shown for unit joint velocity, while at right they are shown for unit joint effort, as calculated according to the power metric. Note that in the plots, each ellipse is at a shape α but is plotted in the space of body velocities ξ , as indicated in the key at the top center, with ξ_θ represented as contour lines on each ellipse.

which we are concerned, suggesting that a more physically-grounded reference choice, such as center of mass and mean orientation or principle moments of inertia, better captures the system maneuverability.

C. Local Maneuverability at Low Reynolds Number

The local maneuverability ellipsoids for the center of mass and mean orientation of the three-link low Reynolds number swimmer are plotted in Fig. 4. Because the swimmer has only two joints, these ellipsoids are flattened into two-dimensional ellipses oriented within a three-dimensional space, as discussed above. At left, the unit shape velocity used to generate the ellipses is taken as unit joint velocity, while at the right, the ellipses are presented for unit effort with respect to the power metric discussed above.

In both cases, the ellipses make it clear that for $\alpha_1 = \alpha_2$ (even, or “C” shapes), the system can move with either pure ξ_y translation or in a coupled (ξ_x, ξ_θ) direction; along the odd axis with $\alpha_1 = -\alpha_2$ (“S” shapes), the system has the options of pure rotation or translation along one (coupled) direction. Between axes, the behavior falls somewhere between these two descriptions, with very narrow ellipses indicating a strong preferred coupling between the translational directions.

The general effect of the change from unit joint velocity to unit joint effort is a decrease in the aspect ratio of the ellipses, indicating that the higher-maneuverability directions cost more than is indicated by using only the joint velocities. One exception to this is along the even axis, where the ξ_x component (which coincides with the minor axis for some, but not all, ellipses) is stretched.

V. CYCLIC MANEUVERABILITY

Maneuverability over large distances depends on the existence of gaits that efficiently translate or rotate the system over each cycle, which we term *cyclic maneuverability*. It is related to local maneuverability – an efficient cycle must have efficient components – but also requires that the segments of the gait not cancel each other’s effects. Gait cycles




Fig. 5: Building the xy cyclic maneuverability polygon. (a) The average velocity of the fastest gait for the system defines both the primary axis of the polygon and a limiting envelope on all other gaits for the system. (b) The optimal-normal gait provides a second limiting envelope. (c) The polygon generated from the optimal and optimal-normal gaits is a lower bound for the cyclic maneuverability. (d) Incorporating more gaits increases this lower bound, especially for directions in which it was originally small.

themselves are important because they represent repeatable shape trajectories that avoid joint limits; controllers built on gaits do not themselves have to exert effort to avoid such constraints.

The optimal gaits discussed in §II-B and §III represent a major component of the cyclic maneuverability, the maximum speed achievable by the system. This speed, however, is for motion in a single, preferred direction. To fully characterize the cyclic maneuverability, we must also consider how easily the system moves orthogonally to this preferred direction, how rapidly it can turn to align the preferred direction with a desired direction of travel, and the turning radius when changing direction – *i.e.*, we need a gait-level equivalent of the maneuverability ellipsoid.

In adapting the maneuverability ellipsoid for use with gaits, three fundamental questions must be answered: What set of gaits should be used as a basis to generate the “ellipsoid”? How should the boundary of the “ellipsoid” be interpolated between points for which it is explicitly calculated? How should these gaits be normalized with respect to each other?

A. The Cyclic Maneuverability Polygon

The finite-dimensionality of the space of joint velocities, together with the linearity of the local connection, makes calculating the local maneuverability ellipsoids relatively straightforward, with the columns of the local connection (the Jacobian) fully-defining their loci. In contrast, the space of gaits is infinite-dimensional, with a nonlinear map from gaits to the net displacements they produce. Even when the gait speeds are normalized, calculating the full set of achievable position velocities would require numerically optimizing a gait to move in each direction of the position space, a clearly infeasible approach, even with the aid of the curvature-based gait evaluation tools from §II-B.

We can, however, use a carefully chosen finite set of gaits to generate a discrete approximation (a polygon) of the cyclic maneuverability that serves as a lower bound for its true value, along with an upper bound that indicates the maximum error between the cyclic maneuverability polygon and the true cyclic maneuverability. For the translational component, we start by taking the fastest pure-translation gait (normalized by shape velocity, not cyclic frequency) and aligning the primary axis of the polygon with the average velocity vector for this gait, as shown in Fig. 5(a). This action immediately provides an upper bound on the xy cyclic

maneuverability, as by definition, no other gait will have a greater speed than this first gait, and thus cannot extend beyond the limiting envelope (i) in Fig. 5(a). To generate a second bound on the maneuverability, we find the gait with optimal speed in the component normal to the first gait, ignoring the parallel component. This gait provides the limiting envelope (ii) in Fig. 5(b) as a second upper-bound for the maneuverability: once again by definition, no gait can have a larger normal component than this one.

With a library of two gaits (and their negatives), we can now construct a cyclic maneuverability polygon in the xy plane. The two gaits are linearly independent, allowing the system to translate in any direction within the plane. Its achievable velocity in a given direction is the weighted average velocity of the two gaits; the set of these velocities forms a polygon whose corners are the gaits in the library, as shown in Fig. 5(c).⁴ This polygon acts as a lower bound on the cyclic maneuverability: the system can move *at least* as quickly as indicated by its edges; any new gaits found will either expand the polygon, as in Fig. 5(d), or will be superseded by faster gaits already in the library. Figures 5(c) and (d) also illustrate the way in which the upper and lower bounds direct the search for additional gaits to include in the library: the second-quadrant edge in (c) is much further from the limiting envelope than is the first-quadrant edge, so the second-quadrant gait added in (d) improves the lower-bound much more than could be expected by incorporating another first-quadrant gait.

B. Rotation

In addition to the efficiency and turning radius of rotational gaits, we must also consider the quantization they engender in the system’s orientation. Such quantization is not generally a concern in translation, where net displacement over any gait is rarely more than a fraction of the system’s body length and the scale of motion is often on the scale of multiple body lengths. Net rotations produced by efficient turning gaits, however, can be large enough to cause problems with precision orientation control. Therefore, in addition to finding the most efficient turning gait, we also, as in the example below, find a range of gaits that are most efficient for turning in smaller increments.

C. Cyclic Maneuverability for the Low Reynolds Number Swimmer

Two cyclic maneuverability polygons for the low Reynolds number swimmer are shown in Fig. 6(a). These polygons correspond to the strokes available to the swimmer under the constraints of no self-intersection (avoiding the kite-shaped regions in the corners of Fig. 6(b), (c), and (d)) or, more restrictively,⁵ $(|\alpha_1|, |\alpha_2|) < 2$. As discussed in §§II–III, speed and efficiency are equivalent at low Reynolds numbers, and the efficiency of a stroke is the ratio of the area integral it encloses on the constraint curvature functions to its (power-normalized) arclength. These properties provide intuition for

⁴As a comparison: the unit shape velocities (interpolated with a 2-norm) in local maneuverability produce an ellipse, while the weighted average (a 1-norm interpolation) in cyclic maneuverability produces a polygon.

⁵*E.g.*, to avoid joint limits or unmodeled fluid coupling between the links.




Fig. 6: Cyclic maneuverability for the low Reynolds number swimmer. Kite-shaped regions in the shape space are self-intersecting configurations.

the construction of the polygons, and seed the numerical optimizations that provide their exact values.

First, the primary axes of both polygons are generated by the optimally-efficient stroke found in [16]. This stroke moves the system along its x axis, as indicated by its strongly sign-definite encirclement of the x component of the curvature function (Fig. 6(b)) and net-zero regions on the y and θ components ((c) and (d)). The lateral bounds on the polygon are then based on the y -efficiency of the strokes, and given either by the second- and fourth-quadrant loops in (c) (the diagonal vectors in (a)), or, if larger angles are allowed, the figure-eight strokes, which produce the pure- y vectors in (a). Whether or not the figure-eights are included, the relative magnitudes of the x and y strokes show that the system has reasonable omnidirectional maneuverability, but is faster when moving in its preferred direction.

The optimal turning strokes in Fig. 6(d) enclose sign-definite regions on the θ curvature function, as deep into the corners as allowed by the constraints. As indicated in Fig. 6(e), the efficiencies (in radians per unit effort) initially grow rapidly with increased displacement per cycle (as the strokes are better able to take advantage of area-perimeter scaling), but then taper off, as the strokes include lower-yield regions of the θ curvature function.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a new framework for assessing the maneuverability of kinematic locomoting systems, along with some key insights regarding motion at low Reynolds numbers. The maneuverability framework is based on earlier results on the manipulability of robotic arms, but takes into account unique aspects of locomotion, such as the role of gaits and selection rules for incorporating new gaits.

As we take this work further, we anticipate its immediate applicability in comparing locomoting systems, both between morphologies, as in [21], and across different media, such as high Reynolds number [11], [12] or granular [14] fluids. We have already incorporated elements of the cyclic maneuverability framework into experiments with the granular swimmers and with our own family of snake robots [25].

We thank the NSF GFRP and Grant 1000389, DARPA RCTA and the Battelle Memorial Institute for supporting this work.

REFERENCES

- [1] S. Kelly and R. M. Murray, "Geometric Phases and Robotic Locomotion," *J. Robotic Systems*, vol. 12, no. 6, pp. 417–431, Jan 1995.
- [2] A. M. Bloch et al., *Nonholonomic Mechanics and Control*. Springer, 2003.
- [3] 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.
- [4] E. A. Shamma, 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.
- [5] R. L. Hatton and H. Choset, "Optimizing coordinate choice for locomoting systems," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, USA, May 2010.
- [6] —, "Connection vector fields for underactuated systems," in *Proceedings of the IEEE BioRobotics Conference*, October 2008, pp. 451–456.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] S. D. Kelly, "The mechanics and control of driftless swimming," in press.
- [11] 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.
- [12] 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.
- [13] E. Kanso, "Swimming Due to Transverse Shape Deformations," *Journal of Fluid Mechanics*, vol. 631, pp. 127–148, 2009.
- [14] 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.
- [15] 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*, July 2011.
- [16] 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.
- [17] A. Shapere and F. Wilczek, "Efficiencies of self-propulsion at low Reynolds number," *J. Fluid Mechanics*, vol. 198, pp. 587–559, 1989.
- [18] —, "Geometry of self-propulsion at low Reynolds number," *Journal of Fluid Mechanics*, vol. 198, pp. 557–585, 1989.
- [19] E. M. Purcell, "Life at low Reynolds numbers," *American Journal of Physics*, vol. 45, no. 1, pp. 3–11, January 1977.
- [20] 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.
- [21] R. L. Hatton and H. Choset, "Kinematic cartography for locomotion," in *Proceedings of Robotics: Science and Systems VII*, Los Angeles, CA USA, June 2011.
- [22] T. Yoshikawa, "Manipulability of Robotic Mechanisms," *The International Journal of Robotics Research*, vol. 4, no. 3, 1985.
- [23] C. M. Gosselin, "The optimum design of robotic manipulators using dexterity indices," *Robotics and Autonomous Systems*, vol. 9, pp. 213–226, 1992.
- [24] G. Pong and J. A. Carretero, "Formulating Jacobian matrices for the dexterity analysis of parallel manipulators," *Mechanism and Machine Theory*, 2006.
- [25] M. Tesch, K. Lipkin, I. Brown, R. Hatton, A. Peck, J. Rembisz, and H. Choset, "Parameterized and Scripted Gaits for Modular Snake Robots," *Advanced Robotics*, vol. 23, no. 9, pp. 1131–1158, 2009.