

Optimal Gaits for Drag-dominated Swimmers with Passive Elastic Joints

Suresh Ramasamy* and Ross L. Hatton†

*Collaborative Robotics and Intelligent Systems (CoRIS) Institute,
Oregon State University, Corvallis, OR, USA*

(Dated: December 7, 2020)

In this paper, we identify optimal swimming strategies for drag-dominated swimmers with a passive elastic joint. We use resistive force theory (RFT) to obtain the dynamics of the system. We then use frequency domain analysis to relate the motion of the passive joint to the motion of the actuated joint. We couple this analysis with elements of the geometric framework introduced in our previous work aimed at identifying useful gaits for systems in drag dominated environments, to identify speed-maximizing and efficiency-maximizing gaits for drag-dominated swimmers with a passive elastic joint.

I. INTRODUCTION

A common strategy for locomotion among animals and robots is to couple cyclical shape changes (gaits) to an interaction with the environment. A long term research focus of the geometric mechanics community has been finding geometric principles that describe what makes a gait effective. Many geometric tools resulting from this line of research have only been developed for systems with a fully actuated shape space. Passive elastic joints, such as flexible fins and tails, however, play an important role in locomotion in many systems [1]. In this paper, we expand on the geometric framework that was developed in [2, 3] to identify optimal gaits for fully actuated swimmers in low Reynolds number fluids, extending it to swimmers with passive elastic shape elements. We then use this extended framework to identify optimal gaits for swimmers in drag dominated environments with passive elastic shape elements.

Our analysis draws on a rich history of work in the geometric mechanics community aimed at using concepts from differential geometry to understand how systems locomote in a low Reynolds number fluid, what the main factors that determine the efficiency of gaits are, and what optimal gaits look like for these systems. In particular, we know that the efficiency of a gait (speed at a given power level or power required at a given speed) depends on its path, period, and pacing in the system's shape space:

1. The net displacement per cycle corresponds to the amount of “constraint curvature” the gait encompasses. Gaits that maximize displacement per cycle thus enclose sign-definite regions of the system's constraint curvature functions (CCFs), which can be obtained from the dynamics of the system [4, 5].

2. At a given level of average power consumption, the minimum time taken to execute one cycle of the gait corresponds to the path length of the gait trajectory under a Riemannian metric on the shape space, meaning that shorter gaits can be executed more frequently or at a lower power level than longer gaits [6].
3. Gaits that progress through their gait kinematics at a steady pace have a lower average power cost for a given frequency than those that “surge” and “dwell” at points in the cycle, or equivalently, can be executed at a higher frequency for a given average power cost [7].

The process of finding efficient gaits thus involves striking a balance between maximizing the enclosed constraint curvature and minimizing the metric-weighted perimeter of the gait and then finding a steady-pace parametrization of the resulting curve. In [2, 3, 8], we built a variational framework based on this geometric insight that identifies optimal gaits for fully-actuated drag-dominated kinematic systems.

In this paper, we extend this framework to identify optimal gaits for swimmers with passive elastic joints and demonstrate the framework on systems with a single actuated and a single passive joint. The ways in which the dynamics of swimmers with passive elastic joints differ from the dynamics of fully actuated swimmers are:

1. Because of the coupling between the actuated and passive joints, the passive swimmers can execute only a subset of the gait kinematics the fully actuated swimmers can execute.
2. This coupling between the actuated and passive joints also endows each gait achievable by the passive swimmer with a unique pacing. Hence, the pacing cost cannot be minimized separately from the kinematic cost.

As illustrated in Fig. 1, we address the two problems introduced by the presence of passive elastic joints: we first use frequency domain analysis to analytically approximate the motion of the unactuated joint in response

* ramasams@oregonstate.edu

† hattonr@oregonstate.edu

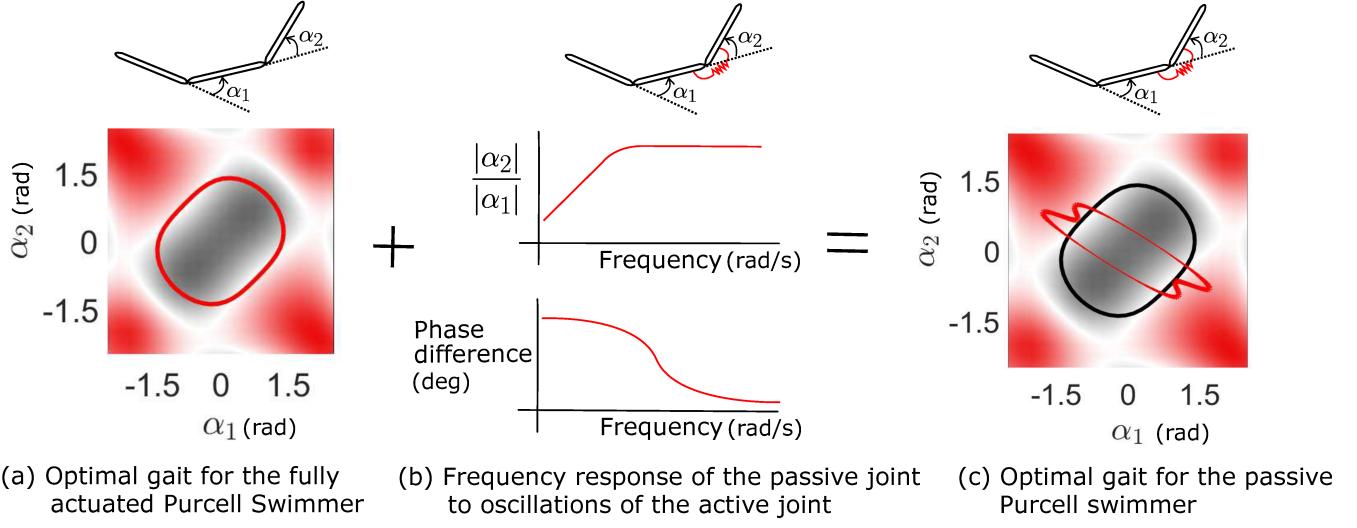


FIG. 1: Effect of a passive elastic joint on the shape and pacing of the (speed-maximizing) optimal gait for a Purcell swimmer. (a) The (speed and efficiency maximizing) optimal gait for the fully actuated Purcell swimmer overlaid on the curvature of the system dynamics. (b) Illustration of the Bode plot of the response of the passive elastic joint to oscillations of the active joint. This response dictates the locus of achievable gaits in the shape space for the passive swimmers. (c) The (speed maximizing) optimal gait for the passive Purcell swimmer. Whereas the optimal gait for the fully actuated Purcell swimmer is plotted with a line of uniform thickness indicating constant power dissipation throughout the cycle, the optimal gait for the passive Purcell swimmer in (c) is plotted using a line of varying thickness, with thickness at a point corresponding to the time spent traversing that section of the gait.

to the motion of the actuated joint. We then combine this frequency-space analysis with elements of the geometric framework introduced in [8] to construct a gradient-descent algorithm that identifies optimal gaits and the pacing associated with these gaits for passive swimmers. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

Drag-dominated swimmers with passive elements were previously studied in [9–13]. Of these works, [9, 10] and [11] are most relevant to this paper, as they discuss the motion of swimmers with a harmonically-driven active joint and a passive joint. The analysis in [9] is particularly relevant, where using perturbation expansion, explicit expressions for leading order solutions were derived for harmonic input oscillations, and the optimal swimmer geometry was obtained for the Purcell swimmer. In this paper, unlike in [9], we do not restrict our input to simple harmonic oscillations, and use a higher order representation of the system dynamics.

The rest of the paper is organized as follows: In §II, we review the geometric locomotion model and derive the dynamics of the drag-dominated swimmers. In §III, we present our frequency domain analysis to relate the motion of the passive joint to the motion of the active joint. In §IV, we recall a few key elements of the gradient descent algorithm introduced in [8], and combine them with the frequency domain analysis of §III, to set up the stroke optimization problem. In §V, we describe the process of finding the speed-maximizing gaits for passive

swimmers. In §VI, we describe the process of finding the efficiency-maximizing gaits for the passive swimmers.

II. SYSTEM DYNAMICS

In this section, we review the dynamics of swimmers in a low Reynolds number fluid. We focus here on the geometric structure of the system dynamics, and not on the explicit expressions. A more detailed discussion of how the dynamics are obtained is presented in [4] and [6]. Some of the the materials presented in this section also appear in [4] and [6].

A. System geometries

The Purcell swimmer illustrated in Fig. 2(a) is a classic example three-link system with minimal complexity that can swim in a low Reynolds number fluid (drag-dominated environment), and was introduced in [14].

The T-link swimmer illustrated in Fig. 2(b) is a modification of the Purcell swimmer where one of the peripheral links is attached to the center link at its midpoint. The T-link swimmer was introduced in [10] as a simplified model to study the swimming dynamics of *Schistosoma mansoni*. *S.mansoni* causes schistosomiasis, a disease comparable to malaria in socio-economic impact [10].

In the passive Purcell swimmer, a torsional spring is attached at the second joint as shown in Fig. 2(a) such that when the joint angle α_2 is zero, the spring is at its equilibrium. In the passive T-link swimmer, the torsional

spring is attached between the middle link and the peripheral link attached at its midpoint as shown in 2(b) such that when the joint angle α_2 is zero, the spring is at its equilibrium.

B. 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 [15].

The locomotion model we employ in this paper was developed for systems in *kinematic* regimes where no gliding can occur i.e., where zero shape velocity results in zero position space velocity [16]. In this model, there exists a linear relationship at each shape between changes in the system's shape and changes in its position,

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

in which $\dot{g} = g^{-1}\dot{g}$ is the body velocity of the system (i.e., \dot{g} expressed in the system's local coordinates), and the *local connection* \mathbf{A} linearly maps joint velocities to the body velocity they produce by pushing the system against its environment. Each row of $-\mathbf{A}$ can be regarded as a body-coordinates local derivative of one position component with respect to the system shape. If we plot the rows of $-\mathbf{A}$ as arrow fields, as in Fig. 3, this means that moving in the direction of the arrows moves the system positively in the corresponding body direction, and moving perpendicular to the arrows results in no motion in that direction [5, 17].

Several efforts in the geometric mechanics community (including our own), have aimed to use the structure of the systems' *Lie brackets* (a measure of how "non-cancelling" the system dynamics are over cyclic inputs) to understand the structure of the optimal solutions to the system equations of motion [4, 6, 18–23]. The core principle in these works is that because the net displacement g_ϕ over a gait cycle ϕ is the line integral of (1) along ϕ , the displacement can be approximated [24] by an integral of the curvature $D(-\mathbf{A})$ of the local connection (its total Lie bracket, which is a measure of how much the connection changes, over a surface ϕ_a bounded by the cycle) [5]:

$$g_\phi = \oint_\phi -g\mathbf{A}(r) \quad (2)$$

$$\approx \exp \iint_{\phi_a} \underbrace{-\mathbf{d}\mathbf{A} + \sum_{i < j} [\mathbf{A}_i, \mathbf{A}_j]}_{D(-\mathbf{A}) \text{ (total Lie bracket)}} , \quad (3)$$

in which $\mathbf{d}\mathbf{A}$ is the exterior derivative of the local con-

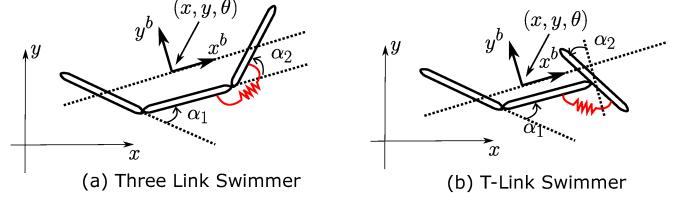


FIG. 2: The geometry and configuration variables of the Purcell swimmer and T-link swimmer

nexion (its generalized row-wise curl),

$$\mathbf{d}\mathbf{A} = \left(\frac{\partial \mathbf{A}_j}{\partial r_i} - \frac{\partial \mathbf{A}_i}{\partial r_j} \right) dr_i \wedge dr_j, \quad (4)$$

which measures how \mathbf{A} changes across the shape space and the local Lie bracket term evaluates (on $SE(2)$) as

$$[\mathbf{A}_i, \mathbf{A}_j] = \begin{bmatrix} \mathbf{A}_i^y \mathbf{A}_j^\theta - \mathbf{A}_j^y \mathbf{A}_i^\theta \\ \mathbf{A}_j^x \mathbf{A}_i^\theta - \mathbf{A}_i^x \mathbf{A}_j^\theta \\ 0 \end{bmatrix} dr_i \wedge dr_j. \quad (5)$$

which measures how forward and turning motions combine into lateral motions via "parallel parking".

In [5, 17, 25], we identified coordinate choices that make the approximation in (3) accurate for large-amplitude gaits, which allowed us to establish a geometric framework to identify and compare gaits that maximized displacement for fully actuated swimmers in drag-dominated environments.

For systems with just two shape variables, plotting the coefficients of the curvature terms as scalar functions on the shape space reveals the attractors that influence the optimal gait cycles: Gaits that produce net displacement in a given (x, y, θ) direction encircle strongly sign-definite regions of the corresponding $D(-\mathbf{A})$ constraint curvature. For example, as illustrated in Fig. 3(b), x -translation gaits encircle the center of the shape space for the Purcell system. In this paper as in [8], we seek gaits that maximize speed and efficiency of motion in the x direction.

Because the scale of motions expected from these systems are generally much larger than displacements produced by executing one gait cycle, the near-optimal way for these systems to move from one point to another is to orient themselves towards the goal and execute gaits that optimize motion in the body x direction. From hereon, g_ϕ will refer to the displacement produced in the x direction by the gait ϕ .

System dynamics using RFT

In this paper, as in [4, 6], we generate the dynamics for our example systems from a resistive force model, in which each element of the body is subject to normal and

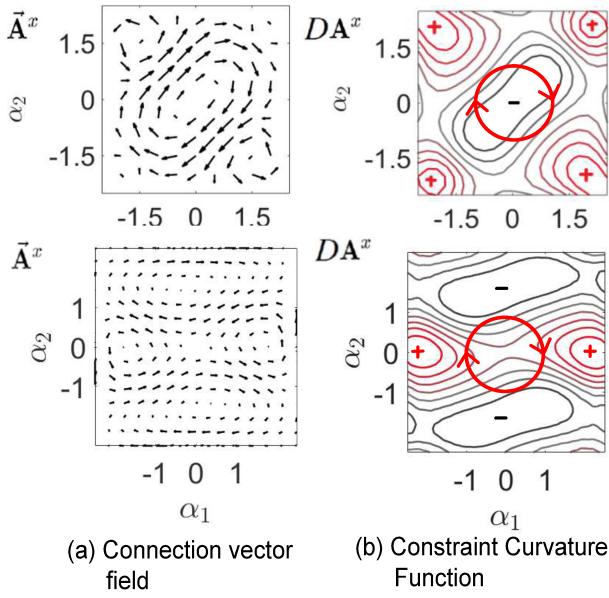


FIG. 3: The connection vector fields in the x -direction for the Purcell swimmer and T-link swimmer are shown in (a) and the corresponding constraint curvature functions are shown in (b). [4].

tangential drag forces proportional to their velocities in those (local) directions [26]. The normal drag coefficient is larger than the tangential component (here, by a factor of 2 : 1), corresponding to the general principle that it is harder to move a slender object in a fluid or on a surface crosswise than it is to move it along its length [27].

We combine this linear drag condition with a quasi-static equilibrium condition that the net drag force and moment on the system is zero at all times (treating the system as heavily overdamped, with acceleration forces much smaller than drag forces). Because the drag forces are not isotropic, the quasi-static condition does not prevent the system from moving, and the system can use the angle-of-attack of its body surfaces to generate net motion.

The quasistatic condition imposes a *Pfaffian constraint* [28] on the system's generalized velocity, which states that combinations of body and shape velocities that are consistent with the system physics must lie in the null space of a linear map from the system velocities to the forces and moments acting on the system's body frame,

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

The matrix ω that maps the velocities to the net forces on the body frame is a pullback of the drag matrices of the individual links via the system kinematics and depends only on the shape r .

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

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

revealing the local connection as $\mathbf{A} = \omega_g^{-1} \omega_r$. The expressions for the dynamics are unwieldy (running to several pages of trigonometric terms in even the simplest cases) so we do not write them out in full here. See [4] for a more detailed treatment of (6)–(7) in the case of the three-link swimmer.

Using the local connection \mathbf{A} and the system's internal kinematics, we can obtain the Jacobian from shape velocity to the local velocity of each section of the body $J(r, \ell)$, where ℓ is the location of the section on the body. We can use this Jacobian to calculate a Riemannian metric \mathcal{M} over the shape space that encodes the cost of effecting a shape change as

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

where C is the matrix of drag coefficients, which acts as a local metric for the motion of each element of the body. For the systems considered in the paper, we take this matrix as,

$$C = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (9)$$

indicating that for any infinitesimal element of a link, the resistance to lateral motion is twice the resistance to longitudinal motion and that the moment-drag on a link is generated by the translation of the longitudinal and lateral forces acting in the local frame of the infinitesimal element to the body frame of the system [29].

As discussed in our previous work [4, 5], the metric \mathcal{M} encodes a quadratic relationship between the shape velocities and power dissipated into the surroundings, given by

$$P = \dot{r}^T \mathcal{M}(r) \dot{r}, \quad (10)$$

as well as the mapping from joint velocities to torques on the joints,

$$\tau = \mathcal{M}(r) \dot{r}. \quad (11)$$

Details of the calculations to generate the local connection \mathbf{A} and the Riemannian metric \mathcal{M} for our example systems are provided in [4] and [6].

In drag-dominated environments, a common measure of the cost of any motion executed by a swimmer is the energy dissipated into the surrounding fluid while executing the motion,

$$E = \int_0^T P(t) dt \quad (12)$$

Substituting the expression for power dissipated from (10) provides a measure of the cost in terms of the metric and shape trajectory,

$$E = \int_0^T P(t)dt = \int_0^T \dot{r}^T \mathcal{M}(r) \dot{r} dt. \quad (13)$$

This cost depends on the geometry, time period and pacing of a given gait. As discussed in [7, 30, 31], for fully actuated swimmers this cost can be broken down into a combination of a pacing-invariant cost that measures the pathlength s of the trajectory through the shape space (weighted by the shapespace metric \mathcal{M}), and a pacing cost σ that measures the deviation from optimal pacing. Finding a gait that minimizes the energy dissipated into the surrounding fluid E is thus equivalent to finding a gait that minimizes the metric-weighted pathlength s , and executing said gait at a constant-power pacing to minimize σ .

The pathlength s of a curve $r(t)$ under a metric \mathcal{M} is

$$s = \oint ds = \oint \sqrt{dr^T \mathcal{M}(r) dr}. \quad (14)$$

Changing the variable of integration from shape r to time t , we can relate the pathlength to the square root of power expended:

$$s = \int_0^T \sqrt{\dot{r}^T \mathcal{M}(r) \dot{r}} dt = \int_0^T \sqrt{P(t)} dt \quad (15)$$

Because moving with constant power is the least-costly pacing with which to execute a motion under viscous drag [7], we can further simplify (15) as

$$s = \sqrt{P_{\text{avg}} T}, \quad (16)$$

where P_{avg} is the average power utilized while executing the motion. This pathlength provides a geometric cost for the best-case execution of the kinematics in a gait cycle.

The additional cost for a non-optimal pacing can be represented by squaring the difference between the average and instantaneous rates at which the gait is being followed (measured as s per time), and then integrating over the time during which the gait is being executed,

$$\sigma = \int_0^{\tau_{\text{total}}} \left(\frac{s_{\text{total}}}{\tau_{\text{total}}} - \frac{d}{d\tau} (s(\tau)) \Big|_{\tau=t} \right)^2 dt, \quad (17)$$

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

Any pacing other than constant-power will make the trajectory take longer for a given average power (or increase the average power required to complete the motion in a fixed time).

III. FREQUENCY DOMAIN ANALYSIS

The key difference in the dynamics of a swimmer with a passive joint when compared to a fully actuated swimmer is the coupling of the motion of the actuated and unactuated joint. In this section, we explore this difference further and present a way of accurately approximating the motion of the unactuated joint from the motion of the actuated joint using frequency domain analysis. The method of linearizing the passive dynamics to obtain approximate limit cycles presented in this section is in the same vein as the limit cycle analysis presented in [13], where a two-link system with static separation between centers of mass and buoyancy was studied.

A. Dynamics of the passive elastic joint

As discussed in [4, 5], the mapping between joint velocities and torques on the joint in the fully actuated swimmers is encoded in the metric calculated in (8),

$$\tau = \mathcal{M}(r)\dot{r}. \quad (18)$$

Because the systems considered in this paper have only one active and one passive joint, this relationship becomes

$$\begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} = \mathcal{M}(\alpha_1, \alpha_2) \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (19)$$

In the case of the swimmers with an elastic joint, because the actuation in the second joint is replaced by an elastic element with stiffness k , the torque τ_2 is always equal to $-k\alpha_2$, i.e.,

$$\begin{bmatrix} \tau_1 \\ -k\alpha_2 \end{bmatrix} = \mathcal{M}(\alpha_1, \alpha_2) \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (20)$$

The first equation in this system of equations,

$$\tau_1 = \mathcal{M}_{11}\dot{\alpha}_1 + \mathcal{M}_{12}\dot{\alpha}_2, \quad (21)$$

thus relates the torque in the actuated joint to the motion of the joints and can be used to calculate the torque required to effect any feasible motion. The second equation in this system of equations,

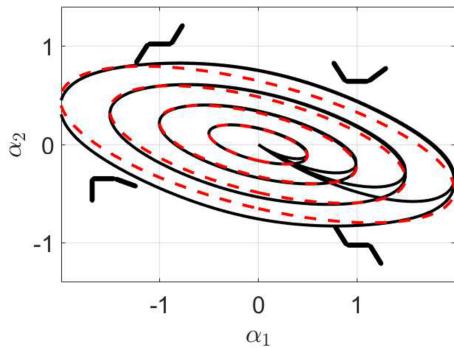
$$-k\alpha_2 - \mathcal{M}_{22}\dot{\alpha}_2 = \mathcal{M}_{21}\dot{\alpha}_1 \quad (22)$$

or equivalently,

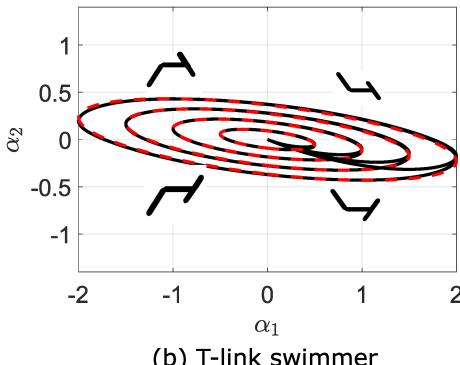
$$-\frac{k}{\mathcal{M}_{21}}\alpha_2 - \frac{\mathcal{M}_{22}}{\mathcal{M}_{21}}\dot{\alpha}_2 = \dot{\alpha}_1 \quad (23)$$

encodes the dynamics of the passive elastic joint in terms of the active joint, and thus defines the space of feasible motions.

Since the joint α_1 is assumed to be the actuated joint, we have full control of $\dot{\alpha}_1$. The value of \mathcal{M} depends on α_1 and α_2 . Its dependence on α_1 and α_2 conveys how



(a) Purcell swimmer



(b) T-link swimmer

FIG. 4: Comparison of the exact stroke limit cycles with the shape predicted by frequency domain analysis. Each black solid line represents the motion of the swimmer when the input to the actuated joint is a sinusoidal wave. We can see the motion converges to a limit cycle. The red dashed line represents the shape of the limit cycles predicted by frequency domain analysis presented in §III. The cartoons in the background show how the Purcell swimmer looks like at different points of the shape space.

the shape of the robot affects the effort required to move the joints. If we consider gaits that are relatively small oscillations of shape, we can approximate the value of \mathcal{M} to be constant throughout the gait. This assumption necessarily introduces errors in our prediction of the motion of the passive joint when the amplitude of input to the active joint is large.

For both the Purcell and T-link swimmers with passive elastic joints, assuming the value of \mathcal{M} to be constant does not introduce significant errors for gaits of amplitude up to 1.5 radians: In Fig. 4, we illustrate the distortions caused in the shape of the limit cycles when we assume \mathcal{M} to be constant throughout the shape space. Each solid black line represents the motion of the full swimmer model when the input to the actuated joint is a sinusoidal wave, and the system starts with both angles at zero. There is a transient term that dominates before the system reaches the limit cycle. The red dashed lines represent the shape of the limit cycles predicted when we assume \mathcal{M} is constant (we describe the calculation of the

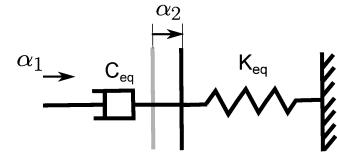


FIG. 5: Spring damper system whose dynamics are equivalent to the passive dynamics of the Purcell swimmer with a passive elastic joint.

limit cycle in the next subsection).

B. Transfer function analysis

In this subsection, our goal is to obtain an analytical approximation of the response of the passive joint to input oscillations of the actively controlled joint. We assume \mathcal{M} to be constant throughout the gait, which makes (23)—the equation that describes the dynamics of the passive elastic joint—a linear first order differential equation and thus well suited to frequency domain analysis.

We use the Laplace transform on (23) to obtain the dynamics of the passive elastic joint in the frequency domain as

$$K_{eq}\mathcal{L}(\alpha_2(t)) + C_{eq}s\mathcal{L}(\alpha_2(t)) = s\mathcal{L}(\alpha_1(t)) \quad (24)$$

where $K_{eq} = \frac{-k}{\mathcal{M}_{21}}$ and $C_{eq} = \frac{-\mathcal{M}_{22}}{\mathcal{M}_{21}}$. These substitutions reveal that the dynamics of the passive joint are equivalent to those of a massless particle attached to a fixed base through a spring and being driven through a damper by a position trajectory α_1 as shown in Fig. 5. Note that the damping coefficient is completely dependent on the physics of the interaction between the swimmer and the fluid.

We can rewrite (24) as a transfer function relation between the active joint α_1 and passive elastic joint α_2

$$\mathcal{L}(\alpha_2(t)) = H(s)\mathcal{L}(\alpha_1(t)) \quad (25)$$

where $H(s) = \frac{s}{(C_{eq}s + K_{eq})}$ is the transfer function that encodes the response of the passive elastic joint to oscillations of the active joint.

Equation (25) tells us how inputs to the controlled joint, α_1 are mapped to the response of the passive joint in the frequency domain. In order to find the response of the passive joint to a sinusoidal oscillation of the actuated joint, we let $\alpha_1(t) = \sin(\omega t)$, then using (25), we

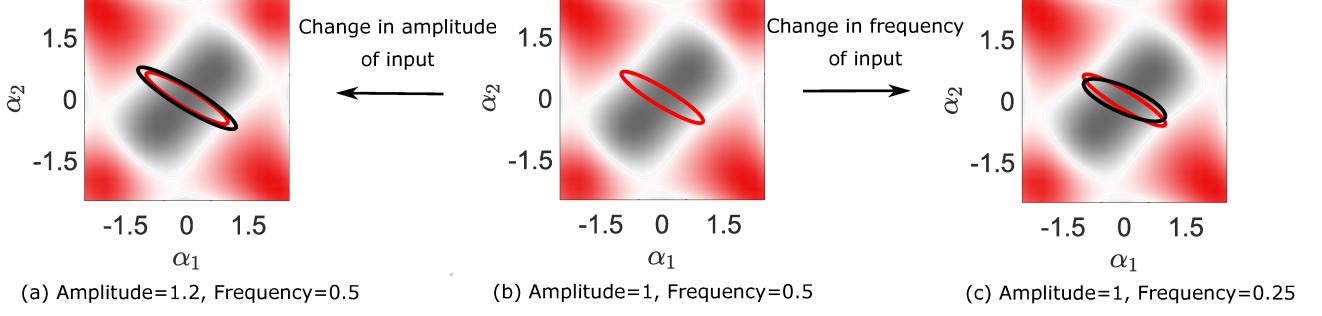


FIG. 6: Changes in the shape of a gait resulting from changes to amplitude and frequency of input oscillation. (a) The gait resulting from a change in amplitude of the nominal input (black) is a scaled version of the gait corresponding to the nominal input actuation (red). (b) The gait resulting from nominal input actuation. (c) The gait resulting from a change in the frequency of the nominal input (black). A change in the frequency of the nominal input leads to a change in both the amplitude and phase of the response of the passive joint.

obtain

$$\mathcal{L}(\alpha_2(t)) = \underbrace{\frac{H(s)}{s}}_{(C_{eq}s + K_{eq})} \underbrace{\frac{\mathcal{L}(\alpha_1(t))}{\omega}}_{(s^2 + \omega^2)} \quad (26)$$

$$= \underbrace{\frac{A_1}{(C_{eq}s + K_{eq})}}_{\text{Transient term}} + \underbrace{\frac{A_2s + A_3}{(s^2 + \omega^2)}}_{\text{Phase shifted sine wave}} \quad (27)$$

where A_1 , A_2 and A_3 can be obtained by equating the two expressions for $\mathcal{L}(\alpha_2(t))$ in (26) and (27) as

$$s\omega = A_1(s^2 + \omega^2) + (A_2s + A_3)(C_{eq}s + K_{eq}), \quad (28)$$

and equating the coefficients of powers of s on each side to extract a system of three equations,

$$\begin{bmatrix} 1 & C_{eq} & 0 \\ 0 & K_{eq} & C_{eq} \\ \omega^2 & 0 & K_{eq} \end{bmatrix} \begin{bmatrix} A_1 \\ A_2 \\ A_3 \end{bmatrix} = \begin{bmatrix} 0 \\ \omega \\ 0 \end{bmatrix}, \quad (29)$$

which can be easily solved to obtain A_1 , A_2 and A_3 as

$$\begin{bmatrix} A_1 \\ A_2 \\ A_3 \end{bmatrix} = \begin{bmatrix} 1 & C_{eq} & 0 \\ 0 & K_{eq} & C_{eq} \\ \omega^2 & 0 & K_{eq} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \omega \\ 0 \end{bmatrix}. \quad (30)$$

If our actuated joint follows a trajectory given by $\alpha_1(t) = B_1 \sin(\omega t)$, after reaching the steady state periodic orbit, the unactuated joint, therefore, follows the trajectory

$$\alpha_2(t) = B_1 A_2 \cos(\omega t) + B_1 A_3 \sin(\omega t). \quad (31)$$

If our actuated joint follows a trajectory given by $\alpha_1(t) = B_1 \sin(\omega t) + B_2 \cos(\omega t)$, after reaching the steady state periodic orbit, the unactuated joint follows the trajectory

$$\begin{aligned} \alpha_2(t) = & B_1 (A_2 \cos(\omega t) + A_3 \sin(\omega t)) \\ & + B_2 (-A_2 \sin(\omega t) + A_3 \cos(\omega t)). \end{aligned} \quad (32)$$

Note that the value of A_2 and A_3 depend on the value of ω , the frequency of the sinusoidal input to the actuated joint. The image of the gait in the shape space is thus coupled to the pacing of the input to the actuated joint.

This coupling between the image of the gait in the shape space and the input to the actuated joint is illustrated in Fig. 6. Fig. 6(b) shows the shape of the gait resulting from a sinusoidal oscillation of the actuated joint of amplitude 1 and frequency 0.5. Fig. 6(a) shows the effect of increasing the amplitude of oscillation on the shape of the gait. As we have linearized the dynamics of the passive joint, a change in frequency produces a scaled version of the original gait. Fig. 6(c) shows the effect of decreasing the frequency of oscillation on the shape of the gait, with a slower oscillation leading to a weaker response (in terms of magnitude) from the passive joint.

The Bode plot of the transfer function H in (25), presented in Fig. 7, illustrates the specific nature of the relationship. We can see that if the frequency of input to the controlled joint is very high, the response of the passive elastic joint is phase shifted relative to the actuated joint by almost 180 degrees, resulting in a gait with a small area. At very low frequencies, the magnitude of the response of the passive elastic joint is small and results in a gait with a small area. At mid-range frequencies, however, the response of the passive elastic joint has a larger amplitude and a phase shift that results in gaits with larger areas.

IV. FINDING OPTIMAL GAITS

Optimal gait design has a long history of research in the physics, mathematics, and engineering communities, as part of the broader field of optimal control [32, 33]. Notable contributions to finding optimal gaits for swimmers in a drag-dominated environment include those of Purcell, who introduced the three-link swimmer as a minimal template for understanding locomotion, a series of

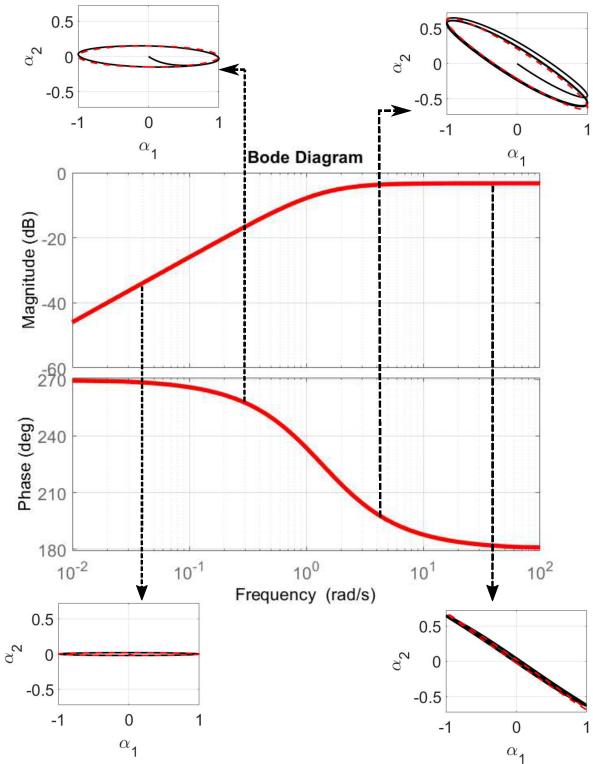


FIG. 7: Bode plot of the transfer function relating the output of the passive joint to the input of the actuated joint in the Purcell swimmer with a passive elastic tail. The inset figures show how the periodic orbits corresponding to gain and phase at certain frequencies look like in the shape space. We can see that actuation at very low frequencies leads to gaits that enclose very little surface area due to the amplitude of the passive joint being low, while at very high frequencies the surface area enclosed by the gait is low due to the almost 180 degree phase shift between the oscillations of the actuated and passive joints.

works [34–38] aimed at numerically optimizing the stroke pattern, and the observation in [7] that the optimal pacing for the gait keeps the power dissipation constant over the cycle.

In drag-dominated environments, a common measure of the cost of any motion executed by a swimmer is the power dissipated into the surrounding fluid while executing the motion. A natural choice of definition for efficiency for these systems is thus

$$\eta_1 = \frac{g_\phi}{E}, \quad (33)$$

where g_ϕ is the displacement produced in the x direction over a gait cycle ϕ and E as defined in (12) is the total energy dissipated into the surrounding while executing the gait ϕ .

Note that this definition of efficiency is the inverse of the mechanical cost of transport used, e.g., in [30]. The mechanical cost of transport is a widely used efficiency metric in the robotics community, especially while studying legged locomotion. As explained in Appendix A, gaits

which optimize our criterion also optimize Lighthill’s efficiency [39], which compares the power dissipated while executing the gait to the power dissipated in rigidly translating the swimmer through the fluid; this measure has an advantage over Lighthill in that it allows for effective comparison between systems with different morphologies and does not require designating a reference shape for rigidly dragging the swimmer.

As noted in §II C, under optimal pacing (constant-power pacing), the pathlength s defined in (14) and the total power dissipated E are equivalent cost criteria. As detailed in [8], for fully actuated swimmers this insight lets us restrict our optimization to constant power trajectories and utilize a geometric definition of efficiency,

$$\eta_2 = \frac{g_\phi}{s} \quad (34)$$

Another important observation to be made here is that, as noted in [2, 3], for the fully actuated swimmers, the gaits that optimized for forward velocity and efficiency were the same within each system. There are two reasons for this:

1. At any given power level P , maximum forward speed is attained by executing the efficiency-maximizing gait at a pacing that has an even power usage instead of surging and sagging.
2. The shape of the efficiency-maximizing gait is invariant across power levels because the displacement resulting from, and the cost incurred by, executing one cycle of any gait at optimal pacing does not depend on the power level P .

The maximum power available, P , dictates the time period T of the speed-maximizing gait. The attainable maximum speed increases linearly with the amount of power available. For efficiency-maximizing gait, the efficiency is invariant across our choice of time period, T , because the displacement resulting from and the cost incurred by executing one cycle of any gait at optimal pacing, does not depend on the power level P . However, the maximum power available P sets a lower bound on the possible range for T .

Moving from an active to a passive swimmer also affects the gradient calculation process that forms the backbone of the framework presented in [8]. In a passive swimmer, the movement of the passive joint is coupled to the movement of the active joint, and therefore we cannot parametrize the motion of the passive joint independently of the motion of the active joint. We discuss the implications of a passive joint on gait parametrization and gradient calculation in §IV B.

A. Efficiency for swimmers with passive elastic joints

In fully actuated swimmers, the transformation of efficiency from the inverse of cost of transport as defined

in (33) to a more geometric definition in (34) was possible because we knew that the optimal pacing keeps the rate of power dissipation constant [7]. In the case of swimmers with passive joints, the response of the passive joint is dictated by the dynamics of the active joint, as illustrated by the Bode plots shown in Fig. 7. As a result, there is a unique pacing associated with every gait the passive swimmer can execute. Changing the pacing of the actuated joint changes the response of the passive joint as shown in §III and hence the shape of the gait. Thus to find the most efficient gait for the Purcell swimmer with a passive elastic joint, we have to directly use the definition of efficiency in (33).

While we could choose a constant power pacing for all gaits for fully actuated swimmers, in the case of the Purcell and T-link swimmers with a passive elastic joint, every gait that respects the passive dynamics of the elastic joint comes with an inherent pacing. Thus for a given spring stiffness of the passive joint, the gait that maximizes forward velocity comes with a power requirement associated with it. Even if we are capable of giving the system more power, there is no way for the system to utilize that power to go faster. Hence for the swimming systems considered in this paper, there are two meaningful measures for comparing different gaits that lead to different definitions of gait optimality: Gaits can be compared by

1. Comparing the average speeds they produce ($\eta = \frac{g\phi}{T}$)
2. Comparing their energetic efficiency (η_1 in (33))

B. Gait parametrization for passive swimmers

We use a truncated Fourier series to parametrize the gaits. This choice of parametrization lets us accurately approximate a large family of smooth periodic gaits. The framework introduced in [8] uses a gradient descent algorithm to identify gaits that maximize efficiency as defined in (34). During the gradient calculation process outlined in appendix B, it is useful to think of the gait as being parametrized by a series of waypoints. We can generate these waypoints from the Fourier parametrization. We use the gradients calculated at each of these waypoints to calculate gradients with respect to the Fourier series parametrization.

In the case of swimmers with a passive joint, we let the actuated joint trajectory $\alpha_1(t)$ be given by a fourth order Fourier series,

$$\alpha_1(t) = a_0 + \sum_{i=1}^4 a_i \cos\left(\frac{2\pi i}{T}t\right) + b_i \sin\left(\frac{2\pi i}{T}t\right) \quad (35)$$

Using (25), i.e. the transfer function relating the movement of the active and passive joints, we obtain the re-

sponse of passive joint to $\alpha_1(t)$ as

$$\alpha_2(t) = \mathcal{L}^{-1}(H(s)\mathcal{L}(\alpha_1(t))). \quad (36)$$

Using explicit evaluation of the transfer function from (27) and (29), we can write the steady state response of the passive joint as

$$\alpha_2(t) = \sum_{i=1}^4 c_i \cos\left(\frac{2\pi i}{T}t\right) + d_i \sin\left(\frac{2\pi i}{T}t\right), \quad (37)$$

where c_i and d_i are functions of a_i and b_i and T .

Using this low-order Fourier series parameterization of the gait, we can generate the direct transcription waypoints, calculate the gradient of the objective function at each waypoint (details of this gradient calculation process for speed-maximizing and efficiency-maximizing gaits are presented in §V and §VI respectively), then project these gradients onto the Fourier basis, to obtain gradients with respect to the Fourier series parameters.

If the system were fully actuated, we could move the Fourier series parameters along these calculated gradient directions to obtain the optimal gait. In the case of passive swimmers, the Fourier coefficients of the unactuated shape direction (c_i, d_i) are functions of the Fourier shape coefficients of actuated shape direction (a_i, b_i). Therefore, to find the correct gradient directions for the Fourier coefficients of the actuated shape, we have to account for the change in the unactuated shape direction that a change in the actuated shape direction would produce.

For an objective function f that depends on a_i, b_i, c_i and d_i , we can calculate the total derivatives of f with respect to a_i and b_i as

$$\frac{df}{da_i} = \frac{\partial f}{\partial a_i} + \frac{\partial c_i}{\partial a_i} \frac{\partial f}{\partial c_i} + \frac{\partial d_i}{\partial a_i} \frac{\partial f}{\partial d_i} \quad (38)$$

$$\frac{df}{db_i} = \frac{\partial f}{\partial b_i} + \frac{\partial c_i}{\partial b_i} \frac{\partial f}{\partial c_i} + \frac{\partial d_i}{\partial b_i} \frac{\partial f}{\partial d_i}, \quad (39)$$

where $\frac{\partial c_i}{\partial a_i}, \frac{\partial c_i}{\partial b_i}, \frac{\partial d_i}{\partial a_i}$ and $\frac{\partial d_i}{\partial b_i}$ are directly taken from the transfer function coefficients (32)

$$\frac{\partial c_i}{\partial a_i} = A_2 \quad (40)$$

$$\frac{\partial c_i}{\partial b_i} = A_3 \quad (41)$$

$$\frac{\partial d_i}{\partial a_i} = -A_2 \quad (42)$$

$$\frac{\partial d_i}{\partial b_i} = A_3. \quad (43)$$

We use these total derivatives to calculate the correct gradient directions for the Fourier coefficients of the actuated shape variables, which account for the fact that in passive swimmers, a change in the shape of input to the actuated shape variable affects the response of the passive elastic joint.

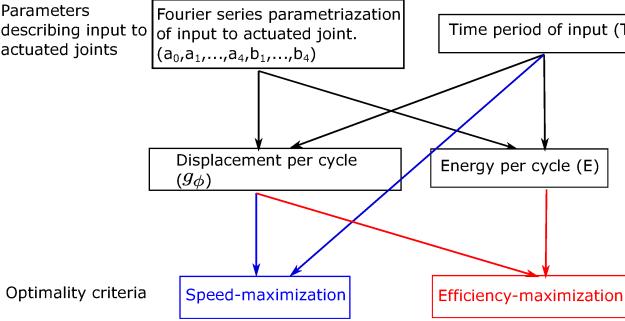


FIG. 8: A flowchart describing how the parameters describing our input to the actuated joint affect the two optimality criteria in the case of drag-dominated swimmers with a passive joint.

V. SPEED-MAXIMIZING GAITS

A major goal of this paper is to geometrically identify gaits for the Purcell and T-link swimmers with an elastic joint that will give it the maximum forward velocity. Therefore the objective function we set out to maximize is $\frac{g_\phi}{T}$, where g_ϕ is the displacement in the body x direction produced on executing the gait ϕ , and T is the time period required to execute the gait.

As explained in §IV.B, we parametrize the actuated shape variable using a low-order Fourier series and obtain the Fourier series parametrization of the resulting motion of the passive joint using the dynamics of the passive elastic joint presented in §III. With these Fourier-series parameters, p_f , we can obtain a sequence of waypoints p_i (in our implementation we use 100 waypoints), equally spaced in time, that describe the location of the discretization points in the shape space. As illustrated in this section, we can then calculate the gradient of speed with respect to each of these waypoints, i.e., calculate the effect moving the waypoints would have on the forward speed attained by the swimmer on executing the gait. We then project these gradients onto the Fourier series basis as explained in §IV.B to obtain the speed-maximizing gaits using gradient descent.

We start from the basic variational principle that functions reach their extrema when their derivatives go to zero. Given a gait parametrization defined by waypoint parameters p , executed in time T , the maximum-velocity cycle must satisfy the condition that the gradient of forward velocity with respect to the parameters p and T is zero, i.e.

$$\frac{\partial}{\partial p} \left(\frac{g_\phi}{T} \right) = \frac{1}{T} \frac{\partial g_\phi}{\partial p} - \frac{g_\phi}{T^2} \frac{\partial T}{\partial p} \quad (44)$$

$$= \frac{1}{T} \frac{\partial g_\phi}{\partial p} = 0, \quad (45)$$

and

$$\frac{\partial}{\partial T} \left(\frac{g_\phi}{T} \right) = \frac{1}{T} \frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{T^2} = 0, \quad (46)$$

where $\frac{\partial T}{\partial p}$ term in (44) can be taken as zero since T and p are two independent variables describing the gait, i.e. the time taken to complete a gait T does not depend on the shape and pacing of the input to the actuated joint described by p .

Once we have the gradient of speed with respect to a direct transcription parametrization p obtained from a Fourier series parametrization p_f , we follow the process outlined in §IV.B, specifically (38) and (39), to obtain the gradient of speed with respect to the Fourier series parametrization, $\frac{\partial}{\partial p_f} \left(\frac{g_\phi}{T} \right)$. A graphical depiction of how elements of the Fourier series parametrization affect the gait optimization process is shown in Fig. 8.

For suitable seed values p_{f0} and T_0 , the maximum-velocity gaits can thus be obtained by finding the equilibrium of the dynamical system

$$\dot{p}_f = \frac{\partial}{\partial p_f} \left(\frac{g_\phi}{T} \right) \quad (47)$$

$$\dot{T} = \frac{\partial}{\partial T} \left(\frac{g_\phi}{T} \right) = \frac{1}{T} \frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{T^2}. \quad (48)$$

In the geometric framework we introduced in [2], we showed that the process of finding efficient gaits for the fully actuated Purcell swimmer is akin to the dynamics of a soap-bubble in which internal pressure and surface tension combine to determine the shape, size and surface concentration of the soap bubble. This is not the case in swimmers with a passive elastic joint. Instead of two independent processes, the optimization process for finding the fastest gait is more unified where (B3) is the equation that helps obtain the shape of the optimal input to the actuated joint and (48) is the equation that helps obtain the optimal pacing of the input.

A. Shape gradient of the optimal input to the actuated joint.

As discussed in [8], the gradient that affects the shape of the input to the actuated joint, $\frac{dg_\phi}{dp}$, pushes the gait towards maximum displacement cycles. From (3), and the fact that variations in p affect ϕ but not the underlying **DA** structure, we can reduce this term to

$$\frac{\partial g_\phi}{\partial p} \approx \nabla_p \iint_{\phi_a} (-D\mathbf{A}). \quad (49)$$

A powerful geometric principle (the general form of the *Leibniz integral rule* [40]) tells us that *the gradient of an integral with respect to variations of its boundary is equal to the gradient of the boundary with respect to these variations, multiplied by the integrand evaluated along the boundary*, and allows us to rewrite (49) in terms of the

constraint curvature value and how changes in parameter values move the gait's trajectory through the shape space.

Formally, the multiplication of the gradient of the boundary and the integrand evaluated along the boundary is the *interior product* [41] of the boundary gradient with the integrand,

$$\nabla_p \iint_{\phi_a} D(-\mathbf{A}) = \oint_{\phi} (\nabla_p \phi) \lrcorner D(-\mathbf{A}). \quad (50)$$

In systems with just two shape variables, the interior product reduces to a simple multiplication between the outward component of $\nabla_p \phi$ and the scalar magnitude of the Lie bracket,

$$\nabla_p \iint_{\phi_a} (-D\mathbf{A}) = \oint_{\phi} (\nabla_{p\perp} \phi) (-D\mathbf{A}). \quad (51)$$

This gradient calculation is illustrated in Fig. 15. The gradient of the enclosed area with respect to variations in the position of p_i , i.e. $\nabla_{p_i} \phi_a$ in the e_{\parallel} and e_{\perp} directions, is the change in triangle's area as p_i moves. Because the triangle's area is always one half base times height (regardless of its pitch or the ratio of its sidelengths), this gradient evaluates to

$$\nabla_{p_i} \phi_a = [e_{\parallel} \ e_{\perp}] \begin{bmatrix} 0 \\ \ell/2 \end{bmatrix}. \quad (52)$$

Note that this term matches the right-hand side of (51), with only normal motions of the boundary affecting the enclosed area.

B. Frequency gradient of the optimal input to the actuated joint

In the case of the fully actuated Purcell swimmer, the shape of the gait, and therefore the displacement produced by executing the gait, are independent of the time taken to execute the gait. This is not true in the case of the Purcell swimmer with a passive elastic joint.

In this subsection, we examine the gradient that guides the optimizer towards the optimal frequency of input to the actuated joint. When the time period required to execute the gait is changed, the shape of the gait changes due to the coupling between the frequency of input to the actuated joint and the response of the passive joint as described in §III. Changing the time period T thus changes not only the frequency of the gait cycle but also the displacement produced per cycle.

We use the chain rule to calculate this gradient,

$$\frac{\partial}{\partial T} \left(\frac{g_{\phi}}{T} \right) = \frac{1}{T} \frac{\partial g_{\phi}}{\partial T} - \frac{g_{\phi}}{T^2} \quad (53)$$

$$= \frac{1}{T} \left(\frac{\partial g_{\phi}}{\partial \alpha_1} \frac{\partial \alpha_1}{\partial T} + \frac{\partial g_{\phi}}{\partial \alpha_2} \frac{\partial \alpha_2}{\partial T} \right) - \frac{g_{\phi}}{T^2} \quad (54)$$

Because α_1 is the actuated shape variable and the shape of the input actuation is independent of the frequency of actuation, $\frac{\partial \alpha_1}{\partial T}$ reduces to zero. Therefore, the gradient of speed with respect to T reduces to

$$\frac{\partial}{\partial T} \left(\frac{g_{\phi}}{T} \right) = \frac{1}{T} \left(\frac{\partial g_{\phi}}{\partial \alpha_2} \frac{\partial \alpha_2}{\partial T} \right) - \frac{g_{\phi}}{T^2}. \quad (55)$$

The first term of the right hand side of (55) captures the contribution to the velocity of the gait caused by the change in the shape of the gait resulting from a change in T . The second term accounts for the fact that, even without a change in the shape of the gait, an increase in the time required to execute the gait would result in a decrease in the velocity of the gait.

C. Passive Purcell and T-Link swimmers

We implemented the optimizer described in the §V in Matlab by providing (47) and (48) as the gradient for the `fmincon` optimizer using the sqp algorithm. The shape of the gait obtained is illustrated in Fig. 9(a). Fig. 9(b) shows the power input to the actuated joint over the cycle. Fig. 9(c) shows a comparison of the speeds achievable by the Passive and fully actuated Purcell swimmers at different power levels. Fig. 10 shows the same results for a passive T-link swimmer.

The transfer function relating the response of the passive joint to oscillations of the input joint is given by (26), $H(s) = \frac{s}{C_{eq}s + K_{eq}}$. Therefore a change in the value of the spring stiffness does not affect the fundamental shape and nature of the response shown by the Bode plot in Fig. 7, but it does shift the entire bode plot to the left or right along the frequency axis. Thus an increase in spring stiffness shifts the Bode plot to the right, which results in the shape of the speed maximizing gait remaining the same, but the time period required to complete the gait decreases, leading to faster speeds.

VI. ENERGY-EFFICIENT GAITS

In this section, we describe the gradient calculations involved in identifying the gait that maximizes the efficiency of the swimmers. The objective function we set out to maximize is

$$\eta = \frac{g_{\phi}}{E}, \quad (56)$$

where g_{ϕ} is the displacement produced on executing the gait ϕ and E is the total energy expended by the robot executing the gait, i.e.

$$E = P_{avg} T \quad (57)$$

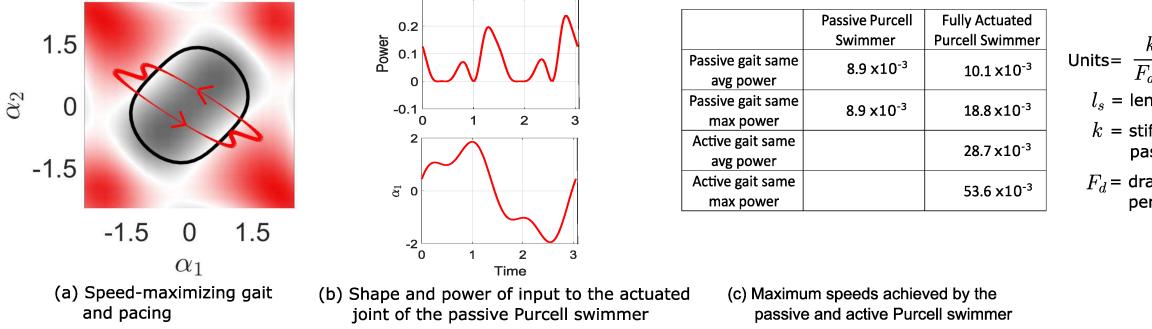


FIG. 9: Gaits that maximize speed along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated Purcell swimmers at different power levels.

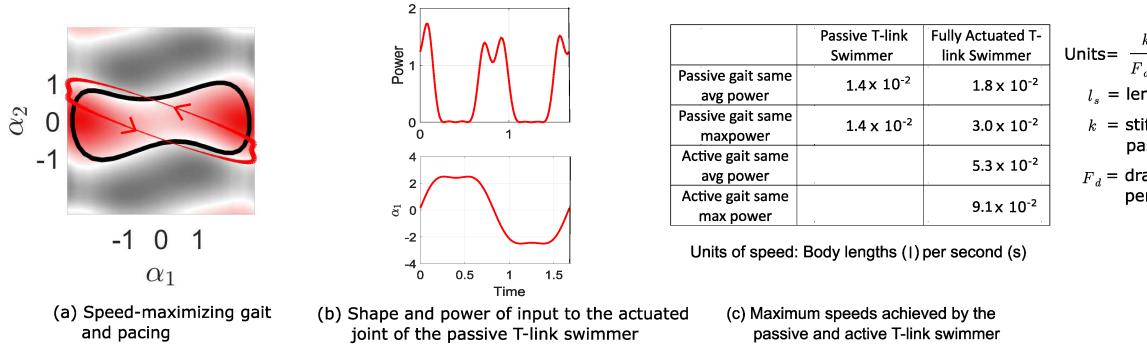


FIG. 10: Gaits that maximize speed along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the speeds achieved by the passive and fully actuated T-link swimmers at different power levels.

As explained in §IV.B, we parametrize the actuated shape variable using a low-order Fourier series, p_f , and obtain the Fourier series parametrization of the resulting motion of the passive joint using the dynamics of the passive elastic joint presented in §III. With these Fourier-series parameters, we can obtain a sequence of waypoints p_i equally spaced in time, that explicitly define the location of the discretization points in the shape space. These waypoints form a direct-transcription parametrization of the gait p , obtained from the Fourier series parametrization p_f . In this section, we present the calculation of the gradient of efficiency at each of these points with respect to p . We then project these gradients onto the Fourier series basis as explained in §IV B to obtain the efficiency maximizing gaits using gradient descent.

The maximum-efficiency cycle must satisfy the condition that the gradient of efficiency with respect to the

parameters, p and T is zero, i.e.

$$\begin{aligned} \frac{\partial}{\partial p} \left(\frac{g_\phi}{E} \right) &= \frac{1}{E} \frac{\partial g_\phi}{\partial p} - \frac{g_\phi}{E^2} \frac{\partial E}{\partial p} \\ &= 0 \end{aligned} \quad (58)$$

and

$$\begin{aligned} \frac{\partial}{\partial T} \left(\frac{g_\phi}{E} \right) &= \frac{1}{E} \frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{E^2} \frac{\partial E}{\partial T} \\ &= 0 \end{aligned} \quad (59)$$

Once we have the gradient of efficiency with respect to a direct transcription parametrization p obtained from a Fourier series parametrization p_f , we follow the process outlined in §IV B to obtain the gradient of efficiency with respect to the Fourier series parametrization, $\frac{\partial}{\partial p_f} \left(\frac{g_\phi}{E} \right)$.

Thus for suitable seed values of p_{f0} and T_0 , the maximum efficiency gaits can be obtained by finding the equilibrium of the dynamical system,

$$\dot{p}_f = \frac{\partial}{\partial p_f} \left(\frac{g_\phi}{E} \right) \quad (60)$$

$$\dot{T} = \frac{1}{E} \frac{\partial g_\phi}{\partial T} - \frac{g_\phi}{E^2} \frac{\partial E}{\partial T} \quad (61)$$

Thus the process of finding the most efficient gait is the result of a unified process where (60) is the equation that helps find the shape of the input to the actuated joint and (61) is the equation that helps find the optimal packing of the input. Note that as in §V, the two equations do not operate independently, and the gradient of shape depends on the time period T , and the gradient of the time period depends on the shape of the gait.

A. Shape gradient of the optimal input to the actuated joint

The shape of the optimal input to the actuated joint is affected by two gradients, $\frac{\partial g_\phi}{\partial p}$ and $\frac{\partial E}{\partial p}$ (60). The details of how $\frac{\partial g_\phi}{\partial p}$ is calculated are explained in §V.A. Whereas $\frac{\partial g_\phi}{\partial p}$ pushes the gait towards maximum displacement cycles, E is a measure of the cost required to execute the gait and $\frac{\partial E}{\partial p}$ pushes the gait towards low cost shapes. At the most efficient gait, these two opposing gradients cancel each other out, and we get an equilibrium for the gait optimization process.

Over a gait cycle, no energy is stored in the spring. Hence we can calculate the energy expended, P , while executing a gait by integrating the power flow through the actuated joint (α_1).

$$E = \int_0^T \dot{\alpha}_1(t)^T \tau_1 dt \quad (62)$$

$$= \int_0^T \dot{\alpha}_1(t)^T \mathcal{M}_1(t) \dot{\alpha}(t) dt \quad (63)$$

where $\mathcal{M}_1(t)$ is the first row of the power metric $\mathcal{M}(t)$. The gradient of cost with respect to the shape of the gait, $\frac{\partial E}{\partial p}$, is calculated by

$$\frac{\partial E}{\partial p} = \frac{\partial}{\partial p} \int_0^T \dot{\alpha}_1(t)^T \mathcal{M}_1(t) \dot{\alpha}(t) dt \quad (64)$$

$$= \int_0^T \left(\frac{\partial \dot{\alpha}_1}{\partial p} \mathcal{M}_1 \dot{\alpha} + \right. \quad (65)$$

$$\left. \dot{\alpha}_1^T \mathcal{M}_1 \frac{\partial \dot{\alpha}}{\partial p} + \dot{\alpha}_1^T \frac{\partial \mathcal{M}_1}{\partial p} \dot{\alpha} \right) dt \quad (66)$$

B. Frequency gradient of the optimal input to the actuated joint

The equation that governs the optimization process for finding the time period of the most efficient gait is de-

scribed by (61). The term $\frac{\partial g_\phi}{\partial T}$ is calculated as described in §V.B. The second gradient in the right hand side of (61) is calculated as

$$\frac{dE}{dT} = \frac{\partial E}{\partial \alpha_1} \frac{\partial \alpha_1}{\partial T} + \frac{\partial E}{\partial \alpha_2} \frac{\partial \alpha_2}{\partial T} + \frac{\partial E}{\partial T}. \quad (67)$$

Because α_1 is the actuated shape variable and the shape of the input actuation is independent of the frequency of actuation, $\frac{\partial \alpha_1}{\partial T}$ reduces to zero. Therefore the gradient of energy with respect to period reduces to

$$\frac{dE}{dT} = \frac{\partial E}{\partial \alpha_2} \frac{\partial \alpha_2}{\partial T} + \frac{\partial E}{\partial T}. \quad (68)$$

The first term accounts for the fact that a change in frequency would change the response of the passive joint α_2 resulting in a change in the shape of the gait, and hence a change in the power dissipated while executing the gait. The second term accounts for the fact that even if the shape of the gait remains unchanged, a change in the frequency of input to the actuated joint will change the time required to execute the gait and hence would change the power dissipated while executing the gait.

C. Passive Purcell and T-Link swimmers

We implemented the optimizer described in the §VI in Matlab by providing the gradients of efficiency with respect to shape and time period, calculated using (60) and (61) respectively, to the `fmincon` optimizer using the sqp algorithm. The shape of the gait obtained is illustrated in Fig. 11(a) for a Purcell swimmer. Fig. 11(b) shows the power input to the actuated joint over the cycle. Fig. 9(c) shows a comparison of the efficiencies achievable by the Passive and fully actuated Purcell swimmers when the forward speed for all the systems is fixed to be equal to the forward speed achieved by the passive swimmer when executing its maximum efficiency cycle. Fig. 12 shows the same results for a passive T-link swimmer. Note that for the Purcell swimmer, the optimizer stops because reducing the frequency further or making the gait smaller does not provide any meaningful increase in efficiency. This observation is line with the results from [9], where the efficiency was found to asymptotically approach a maximum value as frequency of gait oscillations approached zero. The maximum efficiency gait for the passive T-link swimmer is much larger compared to the maximum efficiency gait of the passive Purcell swimmer because a larger gait helps exploit the presence of two peaks in the constraint curvature function of the T-link swimmer.

As discussed in §V.C, change in spring stiffness does not affect the shape and nature of the response of the passive joint, but shifts the bode plot shown in Fig. 7 to the left or right along the frequency axis. An increase in spring stiffness shifts the bode plot to the right, which results in the shape of the efficiency maximizing gait cycle remaining the same, but the time taken to execute the

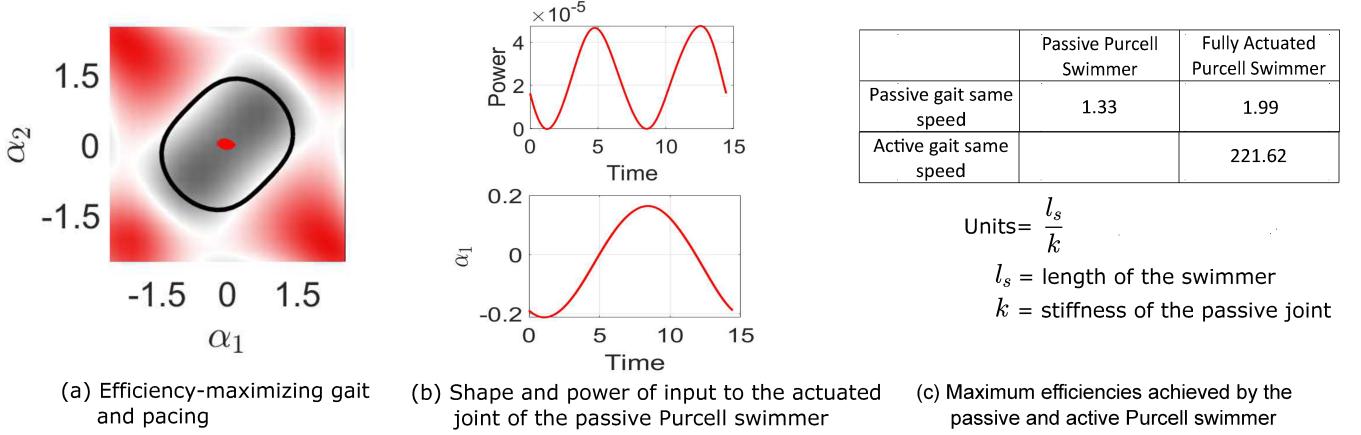


FIG. 11: Gaits that maximize efficiency along x-direction for the Purcell swimmers. (a) The optimal gait for the passive Purcell swimmer (red) and the optimal gait for the fully actuated Purcell swimmer (black). The thickness of the line shows the magnitude of power required at different points of the gait. (b) The power required to execute the optimal gait for the passive Purcell swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated Purcell swimmers moving forward at the same speed as the passive swimmer. Note that for the passive Purcell swimmer, the optimizer stops because reducing the frequency further or making the gait smaller does not provide any meaningful increase in efficiency. This observation is line with the results in [9].

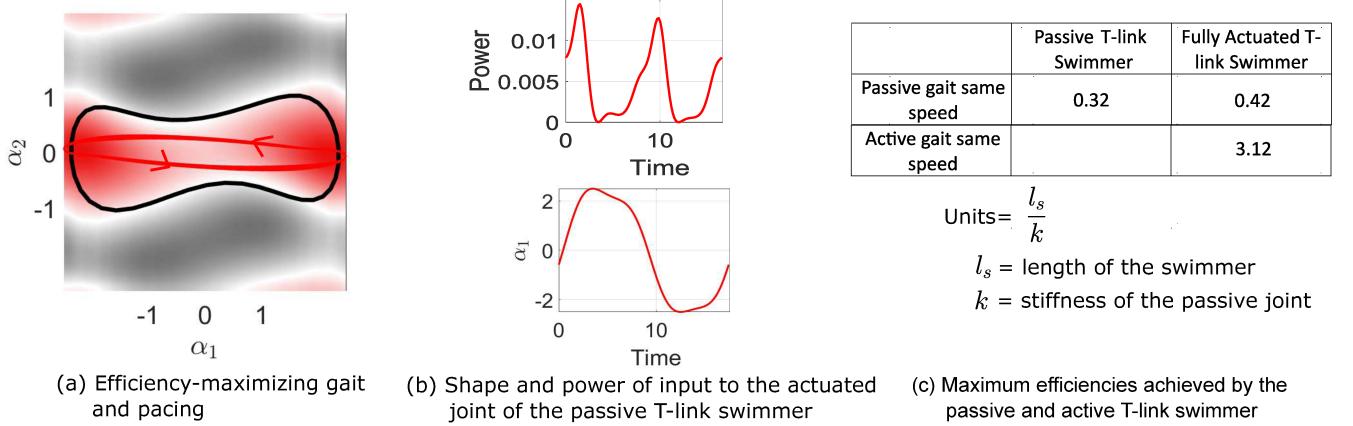


FIG. 12: Gaits that maximize efficiency along x-direction for the T-link swimmers. (a) The optimal gait for the passive T-link swimmer (red) and the optimal gait for the fully actuated T-link swimmer (black). The thickness of the line corresponds to the magnitude of power required at different points of the gait. The maximum efficiency gait of the passive T-link swimmer is much larger compared to the maximum efficiency gait of the passive Purcell swimmer because a larger gait helps exploit the presence of two peaks in the constraint curvature function of the T-link swimmer. (b) The power required to execute the optimal gait for the passive T-link swimmer and input to the actuated joint over one gait cycle. (c) A comparison table of the efficiencies achieved by the passive and fully actuated T-link swimmers moving forward at the same speed as the passive swimmer.

gait will decrease. The energy required to execute a gait is inversely proportional to the time taken to execute the gait. Therefore, more energy is required to execute a gait faster. Thus, increasing spring stiffness will result in an overall decrease in the efficiency of swimming.

VII. COMPARISON WITH PREVIOUS WORK

Passive swimmers have been previously studied, e.g., in [9–13]. Of these works, [10, 11] and [9] are most rele-

vant to this paper, as they discuss the motion of swimmers with a harmonically-driven active joint and a passive joint. The T-link swimmer used in this paper was introduced in [10] as a simplified model to study the swimming dynamics of *Schistosoma mansoni*. The analysis in [9] is particularly relevant, as it applies perturbation theory to investigate the motion of the Purcell swimmer.

The approximation used in the perturbation analysis in [9] is equivalent to assuming that the displacement produced by executing a gait is equal to the area enclosed by the gait in the shape space multiplied by the constraint

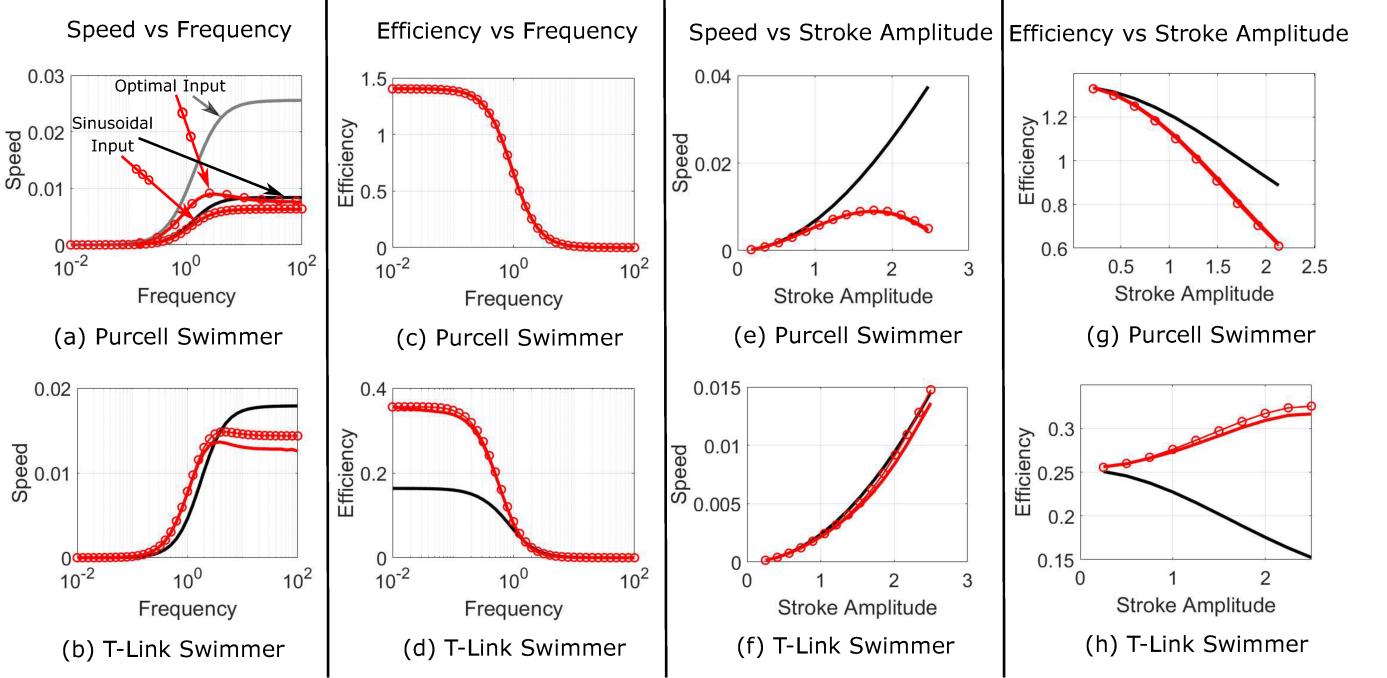


FIG. 13: In all the subfigures, the solid red lines and the red circles show the speeds and efficiencies predicted by numerical simulation and integral of CCF respectively. The solid black lines show the speeds and efficiencies predicted by the constant-CCF assumption used in [9] in the link-attached coordinates. The first column of figures illustrate the speed of the (a) Purcell swimmer as a function of actuation frequency when the input to the controlled joint is a sinusoidal oscillation of unit amplitude and when the input to the controlled joint has the optimal shape and (b) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §V. In (a), the grey line shows the speed predicted by the constant-CCF assumption for a sinusoidal input. The small amplitude perturbation analysis in [9] predicts speed to be a monotonically increasing function of frequency for all inputs. The speed is a monotonically increasing function of frequency for the sinusoidal input. However, it is not monotonically increasing for the optimal input contrary to the prediction from the small amplitude perturbation analysis in [9]. The second column of figures illustrate the efficiency of the (c) Purcell swimmer and (d) T-link swimmer as a function of actuation frequency when the input to the controlled joint has the optimal shape as obtained in §VI. Figures (e)-(h) illustrate the speed and efficiency of the Purcell swimmer and T-link swimmer as a function of the magnitude of optimally shaped input actuation obtained from §VI.

curvature value at the center of the gait,

$$g_\phi \approx \iint_{\phi_a} D(-\mathbf{A})|\mathbf{o} = \phi_a \cdot D(-\mathbf{A})|\mathbf{o}. \quad (69)$$

This approximation has been used in several works from the geometric mechanics community to identify useful shape oscillations that resulted in useful net displacements e.g., [42, 43].

Using the approximation in (69), the authors of [9] concluded that for harmonic inputs, the speed of the swimmer monotonically increases with frequency and asymptotically approaches a maximum value as the actuation frequency approaches ∞ . They similarly concluded that the efficiency of the swimmer asymptotically approaches a maximum value as the actuation frequency approaches zero.

A drawback of this approximation is its the accuracy falls steeply with increasing gait size, as a larger gait would mean larger variations in the value of CCF inside the region bounded by the gait. Our analysis improves on this approximation in three ways: first, our use of a body-averaged frame instead of the link-attached frame

in [9] “flattens” some of the nonlinearity in the system, expanding the domain of gait amplitude for which the perturbation analysis gives accurate results. Second, our use of the full integral of constraint curvature over the area enclosed by the gait,

$$g_\phi \approx \iint_{\phi_a} D(-\mathbf{A}), \quad (70)$$

absorbs much of the remaining nonlinearity. Third, we use (70) only to calculate gradients, and numerically evaluate the value of g_ϕ at each step of the gradient descent process described in §V and §VI to avoid compounding errors from any residual nonlinearity

Fig. 13(a) and Fig. 13(b) show how the effect on swimming speed from changing the frequency of the input stroke for Purcell and T-link swimmers respectively. The solid red lines and the red circles show the speeds predicted by numerical simulation and area integral of curvature respectively. The solid black and grey lines show the speeds predicted by the constant-CCF assumption used in [9] in the link-attached coordinates for optimally shaped and sinusoidal inputs respectively. The speeds predicted by the constant-CCF assumption are higher

than the actual speeds obtained by numerical simulations, and our integral of CCF is a good approximation of the ground-truth simulation.

We can also see from Fig. 13(a) that the velocity obtained by the Purcell swimmer does not monotonically increase with frequency for all inputs to the actuated joint. This results in the speed-maximizing gaits found in §V having an optimal frequency associated with them, rather than exhibiting a monotonic increase in speed with frequency.

In the case of efficiency-maximizing gaits, for the optimal gait shape obtained in §VI, the efficiency does asymptotically approach a maximum value as shown in Fig. 13(c) and Fig. 13(d) as the frequency approaches zero. However, the value is different from the maximum efficiency predicted by applying the small perturbation analysis from [9] to the T-link swimmer, showing that small perturbation analysis does not completely characterize optimal performance. In the case of the Purcell swimmer, the efficiency-maximizing gait is small enough for the perturbation analysis to yield accurate results.

Figs. 13(e)-(h) illustrate how the constant-CCF assumption can introduce errors in identifying optimal actuation shape. Figs. 13(e)-(f) illustrate the effect on the swimming speed from changing the size of the input stroke (the reference input is the optimal input obtained in §V). Figs. 13(g)-(h) illustrate the effect on the swimming efficiency from changing the size of the input stroke (the reference input is the optimal input obtained in §VI). The solid red lines and the red circles show the speeds and efficiencies predicted by numerical simulation and integral of CCF respectively. The solid black lines show the speeds and efficiencies predicted by the constant-CCF assumption used in [9] in the link-attached coordinates.

We can see that in the case of the Purcell swimmer, the constant-CCF assumption used in [9] incorrectly predicts a monotonic increase in speed with an increase in the amplitude of the input to the actuated joint. In the case of the T-link swimmer, the constant-CCF assumption used in [9] incorrectly predicts an increase in efficiency as we shrink the optimal gait. The efficiency would go down if we shrink the optimal gait, because the CCF value for T-link swimmer is higher at the edges than at the center of the shape space. When we shrink the gait, it loses these regions of high value, leading to a decrease in efficiency, which is not captured by the constant-CCF assumption.

The T-link swimmer was first introduced in [10]. The analysis in this paper agrees with the most relevant results from [10], which are that when the actuated joint is driven by a simple harmonic input:

1. There exists a linear relationship between the speed-maximizing value of spring stiffness and frequency of actuation.
2. The average swimming speed increases monotonically as the amplitude of actuation is increased from $\frac{\pi}{2}$ radians to $\frac{11\pi}{9}$ radians.

In §V.C, we noted that an increase or decrease in spring stiffness shifts the Bode plot of the response of the passive joint to the right or to the left without changing the shape of the Bode plot resulting in frequency of the speed-maximizing input being linearly related to the stiffness of the passive joint.

In Fig. 10(a), we can see that the CCF value for T-link swimmer is higher at the edges than at the center of the shape space. Thus an increase in the amplitude of actuation would enclose more of the high value region leading to an increase in speed.

VIII. CONCLUSIONS

In this paper, we have identified the geometric structure of optimal gaits for viscous swimmers with passive elastic joints by combining the constraint-curvature analysis in [8] with frequency-response models for the steady state motion of driven oscillators. We use this structure to identify both speed-maximizing and efficiency-maximizing gaits. The optimal gaits for passive swimmers maximize the CCF integral relative to perimeter and pacing costs, subject to amplitude and phase constraints of a first order system.

As discussed in §IV, for the fully actuated swimmers, the maximum forward speed achievable is only restricted by the maximum power we are able to supply the joints, but for the swimmers with the passive elastic joint, even with more powerful actuators, there is a theoretical maximum forward speed the system can achieve dictated by the stiffness of the passive joint.

The important factor that makes the performance of the fully actuated swimmers superior to that of the swimmers with the passive elastic joint in terms of energy efficiency is the fact that not only can the fully actuated swimmers execute a much larger set of gaits, they can execute any gait the passive swimmer can execute at a pacing that is just as good or better than the pacing dictated by the dynamics of the passive joint.

This raises the question of what benefits, e.g., simplicity of construction, does having a passive elastic member give to biological organisms that locomote in a low Reynolds number fluid? Most biological organisms have tails that resemble an elastic filament. The propulsive and flexible dynamics of such filaments have been well studied [1, 44–46]. Artificial microscopic swimmers with elastic filaments have been proposed based on this body of work [47, 48]. An interesting line of future work would involve investigating the tradeoff between elastic element inefficiencies and structural complexity of being fully actuated.

This work is the first step towards expanding the applicability of the geometric framework presented in [2, 3, 8] to systems where underactuated shape parameters play a role in the dynamics of the system. In the case of the passive Purcell swimmer, assuming the torque required to affect a desired shape change did not depend on the cur-

rent shape of the swimmer did not introduce significant errors in the prediction of the limit cycle corresponding to inputs to the actuated joint as shown in Fig. 4. This might not always be the case in all the swimmers we consider. A future line of research would be to improve our frequency domain analysis by using non-linear perturbation theory to obtain more accurate predictions of limit cycles.

Another line of future work would involve studying the shape of optimal gaits in swimmers with more than one passive joint (e.g. a four-link swimmer with two passive joints). The relationship between design choices (e.g. ratio of link lengths, ratio of joint stiffness) and the shape of the optimal gaits would also be an interesting question to answer in systems with more than one passive joint.

Appendix A: Comparison with Lighthill Efficiency

Lighthill's efficiency, defined as the (reciprocal) ratio between the average power consumed by a given stroke and the power that would be required to drag the swimmer at the same average velocity as that produced by the stroke,

$$\eta_{\text{LH}} = \frac{P_{\text{ref}}}{P_{\text{avg}}}, \quad (\text{A1})$$

is a commonly-used measure of swimming performance.

Because the drag force acting on the swimmer is proportional to the velocity of its motion, it is readily shown that for a given swimmer, Lighthill's efficiency is proportional to $\frac{v_{\text{avg}}^2}{P_{\text{avg}}}$ [9]. Ignoring the proportionality factor, we can thus take the Lighthill efficiency as

$$\eta_{\text{LH}} = \frac{v_{\text{avg}}^2}{P_{\text{avg}}}, \quad (\text{A2})$$

where v_{avg} is the average velocity of the swimmer.

In [8], we use a geometric measure of efficiency, η_2 , defined in (34) as the ratio of displacement produced per cycle to the pathlength of the cycle in the shapespace (weighted by the shapespace metric \mathcal{M}),

$$\eta_2 = \frac{g_\phi}{s} \quad (\text{A3})$$

to identify optimal gaits for fully actuated swimmers. In this appendix, we demonstrate that for fully actuated swimmers the gait that maximizes η_2 , maximizes Lighthill's efficiency and vice versa.

Because s is equal to the time integral of the square root of instantaneous power expended, we can rewrite η_2 as

$$\eta_2 = \frac{g_\phi}{\int_0^T \sqrt{P(t)} dt}. \quad (\text{A4})$$

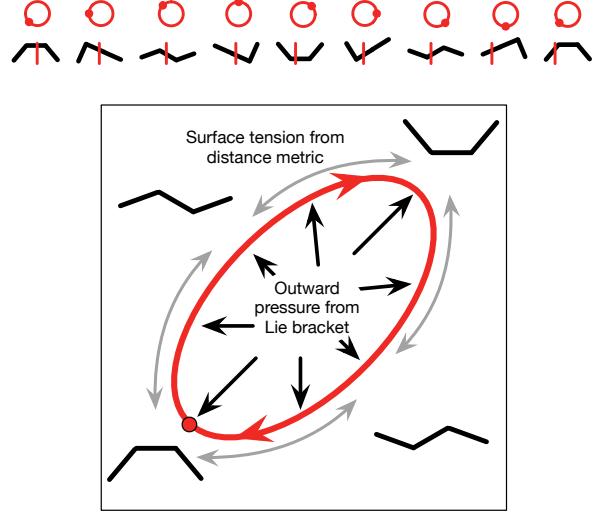


FIG. 14: Our algorithm in [8] maximizes gait efficiency in fully actuated swimmers by finding cycles in the space of body shapes that enclose the most *curvature* of the system dynamics (measured via the Lie bracket) while minimizing their cost-to-execute (measured as the metric-weighted lengths of their perimeters). This process is analogous to the process by which air pressure and surface tension combine to produce the shape and size of a soap bubble. Top: The forward progress of a locomoting system as it executes a gait cycle.

From [7], we know that the optimal pacing for any gait utilizes a constant power pacing and hence for all t ,

$$P(t) = P_{\text{avg}} \quad (\text{A5})$$

Substituting (A5) into the expression for η_2 provides

$$\eta_2 = \frac{g_\phi}{\sqrt{P_{\text{avg}}} T} \quad (\text{A6})$$

$$= \frac{v_{\text{avg}}}{\sqrt{P_{\text{avg}}}} \quad (\text{A7})$$

$$= \sqrt{\eta_{\text{LH}}} \quad (\text{A8})$$

Because the square root is a monotonic function, a gait that maximizes our definition of efficiency with respect to path and is executed at constant power pacing also maximizes the Lighthill efficiency optimized over path and pacing, and vice versa.

Note that this conclusion does not hold for the under-actuated systems we consider in this paper, for which path and pacing cannot be controlled individually.

Appendix B: Soap-bubble Gait Optimization

In this appendix, we present a brief overview of the framework introduced in [8] to identify optimal gaits for fully-actuated drag-dominated swimmers. In §V and §VI, we build on this framework to identify efficiency-maximizing and speed-maximizing gaits for passive

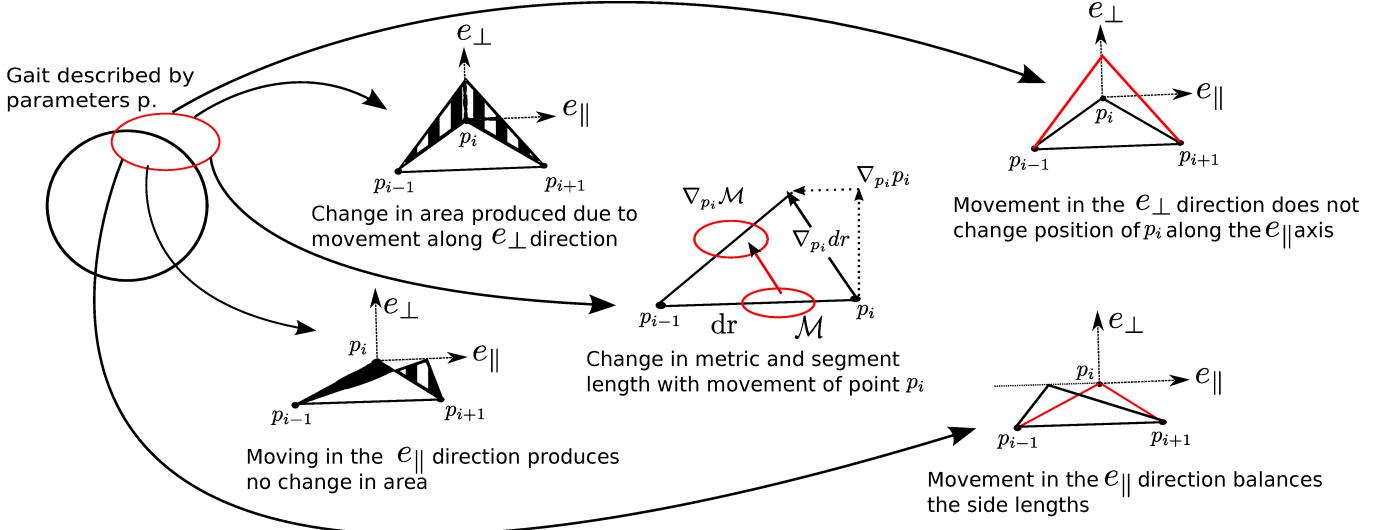


FIG. 15: Changes in area caused by moving in the two coordinate directions in the local frame. Moving in the tangential direction e_{\parallel} produces no change in area, as the area of the triangle given by half the product of base length and height remains the same.

swimmers with one passive and one active joint. The framework uses a gradient descent algorithm to identify gaits that maximize efficiency as defined in (34).

Given a gait parametrization p , maximum-efficiency cycles satisfy the condition that the gradient of the efficiency ratio is zero,

$$\nabla_p \frac{g_\phi}{s} = \frac{1}{s} \nabla_p g_\phi - \frac{g_\phi}{s^2} \nabla_p s = \mathbf{0}. \quad (\text{B1})$$

For suitable seed values p_0 , solutions to (B1) can therefore be reached by finding the equilibrium of the dynamical system

$$\dot{p} = \nabla_p \frac{g_\phi}{s}. \quad (\text{B2})$$

The stable equilibria of the right-hand equation in (B2) are gaits in the same “image families” as the system’s optimally-efficient gaits (i.e., they follow the same curve as the optimal gait, but not necessarily at the same pacing). To construct the optimal gait, we can either optimize via (B2) and then choose a constant-metric-speed parameterization, such that the pacing penalty σ from (17) goes to zero, or directly include $\nabla_p \sigma$ in our optimizer[49].

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

$$\dot{p} = \nabla_p g_\phi - \frac{g_\phi}{s} \nabla_p s + \nabla_p \sigma \quad (\text{B3})$$

(from which we have factored out a coefficient of $\frac{1}{s}$ from (B1)).

As illustrated in Fig. 14, this differential equation is directly analogous to the equations governing the shape of a soap bubble: $\nabla_p g_\phi$ takes the Lie bracket as an “internal

pressure” seeking to expand the gait cycle to fully encircle a sign-definite region, $\nabla_p s$ is the “surface tension” that constrains the growth of the bubble, and $\nabla_p \sigma$ is the “concentration gradient” that spreads the soap over the bubble’s surface.

For the fully actuated swimmers, we parametrize the gait as a sequence of waypoints p_i such that the gait parameters p_i explicitly define the location of the discretization points. As illustrated in Fig. 15, each waypoint p_i forms a triangle with its neighboring points and we can define a local tangent direction e_{\parallel} as

$$p_{i+1} - p_{i-1} = \ell e_{\parallel} \quad (\text{B4})$$

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

We selected this direct-transcription parameterization because it facilitated visualizing the workings of our optimizer (and thus the dynamics governing any other optimization applied to this problem). Additionally, it allowed us to illustrate simultaneous optimization of the gait path and its pacing. We could also parametrize the gait using a Fourier series or Legendre polynomials. In this case, the pacing optimization should be done after the image of the optimal gait has been found because finding an optimal pacing can no longer be formulated as a process orthogonal to the gradient descent process for finding the image of the optimal gait.

ACKNOWLEDGEMENT

This work was supported by the National Science Foundation, under CMMI grants 1462555 and 1653220. The authors would like to thank Shai Revzen and Brian Bittner for several productive discussions that con-

tributed to this work, and Saad Bhamla for bringing the work in [10] to our attention.

[1] C. Brennen and H. Winet, Fluid mechanics of propulsion by cilia and flagella, *Annual Review of Fluid Mechanics* **9**, 339 (1977).

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

[3] S. Ramasamy and R. L. Hatton, Geometric gait optimization beyond two dimensions, in *American Control Conference (ACC), 2017* (IEEE, 2017) pp. 642–648.

[4] R. L. Hatton and H. Choset, Geometric swimming at low and high Reynolds numbers, *IEEE Transactions on Robotics* **29**, 615 (2013).

[5] R. L. Hatton and H. Choset, Nonconservativity and non-commutativity in locomotion, *European Physical Journal Special Topics: Dynamics of Animal Systems* **224**, 3141 (2015).

[6] R. L. Hatton, T. Dear, and H. Choset, Kinematic cartography and the efficiency of viscous swimming, *IEEE Transactions on Robotics* (2017).

[7] 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* **490**, 15 (2003).

[8] S. Ramasamy and R. L. Hatton, The geometry of optimal gaits for drag-dominated kinematic systems, *IEEE Transactions on Robotics* **35**, 1014 (2019).

[9] E. Passov and Y. Or, Dynamics of purcell’s three-link microswimmer with a passive elastic tail, *The European Physical Journal E* **35**, 1 (2012).

[10] D. Krishnamurthy, G. Katsikis, A. Bhargava, and M. Prakash, Schistosoma mansoni cercariae swim efficiently by exploiting an elastohydrodynamic coupling, *Nature Physics* **13**, 266 (2017).

[11] A. Montino and A. DeSimone, Three-sphere low-reynolds-number swimmer with a passive elastic arm, *The European Physical Journal E* **38**, 42 (2015).

[12] I. Jo, Y. Huang, W. Zimmermann, and E. Kanso, Passive swimming in viscous oscillatory flows, *Physical Review E* **94**, 063116 (2016).

[13] L. J. Burton, R. L. Hatton, H. Choset, and A. E. Hosoi, Two-link swimming using buoyant orientation, *Physics of Fluids* **22**, 091703 (2010).

[14] E. M. Purcell, Life at low Reynolds numbers, *American Journal of Physics* **45**, 3 (1977).

[15] In 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*.

[16] This kinematic condition has been demonstrated for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [19, 22], non-holonomic constraints such as passive wheels [20, 22, 42, 50], and fluid interactions at the extremes of low [4, 23] and high [4, 21, 51] Reynolds numbers.

[17] R. L. Hatton and H. Choset, Geometric motion planning: The local connection, Stokes’ theorem, and the importance of coordinate choice, *International Journal of Robotics Research* **30**, 988 (2011).

[18] S. D. K. Kelly and R. M. Murray, Geometric phases and robotic locomotion, *J. Robotic Systems* **12**, 417 (1995).

[19] G. C. Walsh and S. Sastry, On reorienting linked rigid bodies using internal motions, *Robotics and Automation, IEEE Transactions on* **11**, 139 (1995).

[20] J. P. Ostrowski and J. Burdick, The mechanics and control of undulatory locomotion, *International Journal of Robotics Research* **17**, 683 (1998).

[21] J. B. Melli, C. W. Rowley, and D. S. Ruffat, Motion planning for an articulated body in a perfect planar fluid, *SIAM Journal of Applied Dynamical Systems* **5**, 650 (2006).

[22] E. A. Shammas, H. Choset, and A. A. Rizzi, Geometric motion planning analysis for two classes of underactuated mechanical systems, *Int. J. of Robotics Research* **26**, 1043 (2007).

[23] J. E. Avron and O. Raz, A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin, *New Journal of Physics* **9** (2008).

[24] This 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 [52, 53]. For a discussion of the accuracy of this approximation and its derivation, see [5].

[25] 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, 2010) pp. 4493–4498.

[26] This model is most widely associated with swimmers at low Reynolds numbers (e.g., [34]), but can also be regarded as an informative general model for systems that experience more lateral drag than longitudinal drag (e.g. [54]). Our choice of resistive force here also does not preclude the use of more dejointed physical models (e.g., [55]) to construct the local connection \mathbf{A} .

[27] Note that a more complete fluid model for low Reynolds number swimming including interbody flow interactions would change numbers in the dynamics but not the overall structure.

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

[29] This metric corresponds to approximating the viscous drag acting on the swimmer via a local resistive force model. If we used a more accurate fluid model, the drag on a low-Reynolds number swimmer would have the same form, except that C would also depend on r and ℓ .

[30] J. E. Avron, O. Gat, and O. Kenneth, Optimal swimming at low reynolds numbers, *Physical Review Letters* **93**, 186001 (2004).

[31] R. L. Hatton and H. Choset, Kinematic cartography and the efficiency of viscous swimming, *IEEE Transactions on Robotics* **33**, 523 (2017).

[32] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control, Optimization, Estimation, and Control* (John Wiley & Sons, New York-London-Sydney-Toronto, 1975) pp. 402–402.

[33] J. P. Ostrowski, J. P. Desai, and V. Kumar, Optimal gait selection for nonholonomic locomotion systems, *International Journal of Robotics Research* **19**, 225 (2000).

[34] D. Tam and A. E. Hosoi, Optimal stroke patterns for Purcell’s three-link swimmer, *Phys. Review Letters* **98**, 068105 (2007).

[35] A. DeSimone, L. Heltai, F. Alouges, and A. Lefebvre-Lepot, Computing optimal strokes for low Reynolds number swimmers, *Natural Locomotion in Fluids and on Surfaces* , 177 (2012).

[36] L. Giraldi, P. Martinon, and M. Zoppello, Controllability and optimal strokes for n-link microswimmer, in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on* (2013) pp. 3870–3875.

[37] P. Bettoli, B. Bonnard, L. Giraldi, P. Martinon, and J. Rouot, The Purcell Three-link swimmer: some geometric and numerical aspects related to periodic optimal controls (2015), working paper or preprint.

[38] A. Shapere and F. Wilczek, Geometry of self-propulsion at low Reynolds number, *Journal of Fluid Mechanics* **198**, 557 (1989).

[39] J. Lighthill, Flagellar hydrodynamics, *SIAM Review* **18**, 161 (1976).

[40] H. Flanders, Differentiation under the integral sign, *The American Mathematical Monthly* **80**, 615 (1973).

[41] Not the inner product, see [40] for more details.

[42] R. M. Murray and S. S. Sastry, Nonholonomic motion planning: Steering using sinusoids, *IEEE Transactions on Automatic Control* **38**, 700 (1993).

[43] K. A. Morgansen, B. I. Triplett, and D. J. Klein, Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles, *IEEE Transactions on Robotics* **23**, 1184 (2007).

[44] C. H. Wiggins and R. E. Goldstein, Flexive and propulsive dynamics of elastica at low reynolds number, *Physical Review Letters* **80**, 3879 (1998).

[45] R. E. Goldstein and S. A. Langer, Nonlinear dynamics of stiff polymers, *Physical review letters* **75**, 1094 (1995).

[46] T. S. Yu, E. Lauga, and A. Hosoi, Experimental investigations of elastic tail propulsion at low reynolds number, *Physics of Fluids* **18**, 091701 (2006).

[47] R. Dreyfus, J. Baudry, M. L. Roper, M. Fermigier, H. A. Stone, and J. Bibette, Microscopic artificial swimmers, *Nature* **437**, 862 (2005).

[48] J. Edd, S. Payen, B. Rubinsky, M. L. Stoller, and M. Sitti, Biomimetic propulsion for a swimming surgical micro-robot, in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, Vol. 3 (IEEE, 2003) pp. 2583–2588.

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

[50] A. M. Bloch *et al.*, *Nonholonomic Mechanics and Control* (Springer, 2003).

[51] E. Kanso, Swimming due to transverse shape deformations, *Journal of Fluid Mechanics* **631**, 127 (2009).

[52] J. E. Radford and J. W. Burdick, Local motion planning for nonholonomic control systems evolving on principal bundles, in *Proceedings of the International Symposium on Mathematical Theory of Networks and Systems* (Padova, Italy, 1998).

[53] W. Magnus, On the exponential solution of differential equations for a linear operator, *Communications on Pure and Applied Mathematics* **VII**, 649 (1954).

[54] R. L. Hatton, H. Choset, Y. Ding, and D. I. Goldman, Geometric visualization of self-propulsion in a complex medium, *Physical Review Letters* **110**, 078101 (2013).

[55] N. Giuliani, L. Heltai, and A. DeSimone, Predicting and optimizing microswimmer performance from the hydrodynamics of its components: The relevance of interactions, *Soft robotics* **5**, 410 (2018).