

Learning to swim in potential flow

Yusheng Jiao ¹, Feng Ling ¹, Sina Heydari ¹, Nicolas Heess,² Josh Merel,² and Eva Kanso^{1,*}

¹*Department of Aerospace and Mechanical Engineering, University of Southern California,
Los Angeles, California 90089, USA*

²*DeepMind, London, United Kingdom*



(Received 8 August 2020; accepted 15 March 2021; published 12 May 2021;
corrected 13 May 2021)

Fish swim by undulating their bodies. These propulsive motions require coordinated shape changes of a body that interacts with its fluid environment, but the specific shape coordination that leads to robust turning and swimming motions remains unclear. To address the problem of underwater motion planning, we propose a simple model of a three-link fish swimming in a potential flow environment and we use model-free reinforcement learning for shape control. We arrive at optimal shape changes for two swimming tasks: swimming in a desired direction and swimming towards a known target. This fish model belongs to a class of problems in geometric mechanics, known as *driftless* dynamical systems, which allow us to analyze the swimming behavior in terms of geometric phases over the shape space of the fish. These geometric methods are less intuitive in the presence of *drift*. Here, we use the shape space analysis as a tool for assessing, visualizing, and interpreting the control policies obtained via reinforcement learning in the absence of drift. We then examine the robustness of these policies to drift-related perturbations. Although the fish has no direct control over the drift itself, it learns to take advantage of the presence of moderate drift to reach its target.

DOI: [10.1103/PhysRevFluids.6.050505](https://doi.org/10.1103/PhysRevFluids.6.050505)

I. INTRODUCTION

Fish swim through interactions of body deformations with the fluid environment. A fish assimilates sensory information about its own body and the external environment and produces patterns of muscle activation that result in desired body deformations [see Fig. 1(a)]. How these sensorimotor decisions are enacted at the physiological level, at the level of neuronal circuits, remains unclear [1–4]. Animal models, such as the *Danio rerio* zebrafish [5–7], as well as robotic and mathematical models [8,9], provide valuable insights into the sensorimotor control underlying fish behavior. Such understanding offers enticing paradigms for the design of artificial soft robotic systems in which the control is *embodied* in the physics [10,11]. Embodied systems sense and respond to their environment through a physical body [12,13]; physical interactions with the environment are thus vital for sensing and control. In fish, the dynamics of the fluid environment is essential both at an evolutionary time scale—in shaping body morphologies [14] and sensorimotor modalities—and at a behavioral time scale. Fish bodies are tuned to exploit flows [15,16]. Body designs and undulatory motions have been examined in computational and semianalytical models of fluid-structure interactions [17–21], including models of body stiffness and neuromechanical control [22,23]. The fish’s ability to integrate multiple sensory modalities such as vision [24–26] and flow sensing [27,28] are

*kanso@usc.edu; <https://sites.usc.edu/kansolab/>

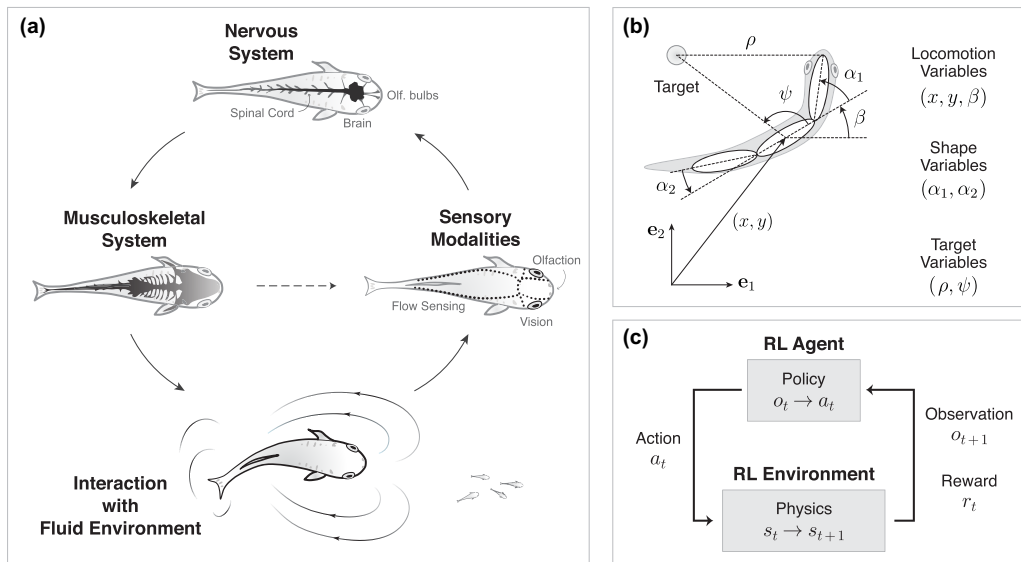


FIG. 1. Model-free reinforcement learning and the three-link fish. (a) Illustration of sensorimotor feedback loops in fish. Motor commands generated in the nervous system activate the musculoskeletal system, resulting in deformations of the body. Body deformations, through interaction with the fluid environment, lead to swimming; meanwhile, sensory modalities provide sensory feedback to the nervous system. The dashed arrow between musculoskeletal and sensory systems indicates somatic sensing used to assess whether previous motor commands were successfully executed. Other reflexive or preflex signals could also be at play [2,27,46]. (b) Three-link fish swimming in quiescent fluid. Locomotion variables (x, y, β) are set in a laboratory fixed frame, while the shape variables (α_1, α_2) and target variables (ρ, ψ) are set in a body frame symbolizing egocentric control and learning. (c) To apply model-free reinforcement learning to our problem, we only need to set the appropriate state, observation, action, and reward variables based on our fish model.

essential to behaviors ranging from rheotaxis [29–31] to schooling [32–35]. Recent developments prove that machine learning techniques are highly effective in addressing problems of flow sensing and fish behavior [36–43].

A central problem in fish behavior, which is also relevant for underwater robotic systems, is *gait design* or *motion planning*: what body deformations produce a desired swimming objective? The answer requires an understanding of how the numerous biomechanical degrees of freedom of the fish body are coordinated to achieve the common objective. Mathematically, this problem is often expressed in terms of an optimality principle: find control laws that optimize a desired objective, such as maximizing swimming speed or minimizing energetic cost [17,18,44]. But how these control laws are implemented in the nervous system, and how they are acquired via learning algorithms, are typically beyond the scope of such methods. As these optimization methods rely on an internal model of the dynamics, different computational results have been obtained by varying the specification of the physical model of the fish, the performance metric, and the control constraints [17,18,44,45].

Model-free reinforcement learning (RL) of embodied systems offers an alternative framework for gait design that is mathematically and computationally tractable [47–49]. In this framework, the fish’s world is divided into a body controlled by a *learning agent* and an environment that encompasses everything outside of what the agent can control. The agent can be viewed as an abstract representation of the parts of the fish responsible for sensorimotor decisions. In RL, the agent must learn from experience in a trial-and-error fashion. Specifically, repeated interactions of the body with the environment enable the agent to collect *sensory observations*, *control actions*,

and associated *rewards*. The goal of the agent is thus to learn to produce behavior that maximizes rewards, and the process is model-free when learning does not make use of either *a priori* or developed knowledge of the physics of the system. The learned feedback control law, called a *policy*, is essentially a mapping from sensory observations to control actions. This mapping is nonlinear and stochastic, and, by construction, rather than providing a single optimal trajectory, it can be applied to any initial condition and transferred to conditions other than those seen during training such as when the body or fluid environment are perturbed [50].

Here, we employ RL to design swimming gaits. We use an idealization in which the fish is modeled as an articulated body consisting of three links, with front and rear links free to rotate relative to the middle link via hinge joints [44,45,51–54]. In describing the physics of the fish, we cede the complexity of accounting for the full details of the fluid medium in favor of considering momentum exchange between the articulated body and the surrounding fluid in the context of a potential flow model [44,45,52,53]. This model is a canonical example of a class of under-actuated control problems whose dynamics can be described over the actuation (shape) space using tools from geometric mechanics [44,55–58]. Specifically, swimming motions can be represented by the sum of a *dynamic phase* or *drift*, and a *geometric phase* over the shape space of all body deformations [45]. In the geometric phase case, a *geometric connection*, defined as vector fields over the actuation shape space, maps infinitesimal shape changes into infinitesimal rigid motions of the whole body. From a motion control perspective, this geometric framework is advantageous in that it provides tools for gait analysis and manual design of control policies over the full shape space [54]. These geometric tools also provide an intuitive way for direct and interpretable visualizations of the RL-based policies. Here, we visualize the RL policies as vector fields over the shape space. This framework leads to a straightforward interpretation of the RL policy: given an observation of its own shape, the fish’s action needs to follow the corresponding vector in the shape space. A trajectory or realization of the RL policy arises from locally following these vectors to achieve a desired global task. These tools also allow us to probe the optimality and behavior of the RL policy in light of the physics of the system and in comparison to manually designed policies over the fish shape space.

II. MATHEMATICAL MODEL OF THE THREE-LINK FISH

Consider a three-link fish as shown in Fig. 1(b). Rotations of the front and rear links relative to the middle link are denoted by the angles α_1 and α_2 such that (α_1, α_2) fully describe all possible body deformations. We constrain the swimming motions to a two-dimensional plane, and let (x, y) and β denote the net planar displacement and rotation of the fish body, such that (\dot{x}, \dot{y}) and $\dot{\beta}$ represent the linear and rotational velocities of the fish in the inertial frame [the dot notation $(\dot{}) = d()/dt$ represents differentiation with respect to time t]. We also introduce the linear velocity (v_1, v_2) expressed in a co-rotating body frame attached to the center of the middle link,

$$v_1 = \dot{x} \cos \beta + \dot{y} \sin \beta, \quad v_2 = -\dot{x} \sin \beta + \dot{y} \cos \beta. \quad (1)$$

The total linear momentum (p_x, p_y) and total angular momentum π of the body-fluid system are expressed in the inertial frame, and they are related to their counterparts (P_1, P_2) and Π in the body-fixed frame as follows:

$$p_x = P_1 \cos \beta - P_2 \sin \beta, \quad p_y = P_1 \sin \beta + P_2 \cos \beta, \quad \pi = \Pi. \quad (2)$$

In potential flow, it is a known result that the total linear and angular momenta of the body-fluid system can be expressed in terms of the body geometry and velocity, via the so-called added mass matrices [44,45,59]. Expressions for the added mass matrices of the three-link fish when the links are hydrodynamically decoupled (coupled only geometrically through the holonomic constraints at

the joints) are derived in detail in Appendix A. The total momenta (P_1, P_2) and Π are given by

$$\begin{bmatrix} P_1 \\ P_2 \\ \Pi \end{bmatrix} = \mathbb{I}_{\text{lock}} \begin{bmatrix} v_1 \\ v_2 \\ \dot{\beta} \end{bmatrix} + \mathbb{I}_{\text{couple}} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}, \quad (3)$$

where \mathbb{I}_{lock} is the locked mass matrix at a given shape of the fish [see Eqs. (A6)–(A9)] and $\mathbb{I}_{\text{couple}}$ is the mass matrix associated with shape deformations [see Eq. (A10)].

In the absence of external forces and moments on the fish-fluid system, the total momenta (p_x, p_y) and π are conserved for all time. Conservation of total momentum yields, upon inverting Eqs. (2) and substituting into Eq. (3),

$$\begin{bmatrix} v_1 \\ v_2 \\ \dot{\beta} \end{bmatrix} = \mathbb{I}_{\text{lock}}^{-1} \begin{bmatrix} \cos \beta & \sin \beta & 0 \\ -\sin \beta & \cos \beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} p_x \\ p_y \\ \pi \end{bmatrix} - \mathbb{I}_{\text{lock}}^{-1} \mathbb{I}_{\text{couple}} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (4)$$

If we further substitute Eqs. (1) into Eq. (4), we arrive at three coupled first-order equations of motion for x, y , and β given α_1 and α_2 . The control problem consists of finding the time evolution of shape changes ($\alpha_1(t), \alpha_2(t)$) that achieve a desired locomotion task ($x(t), y(t)$) and $\beta(t)$. Specifically, a swimming gait is defined as a cyclic shape change ($\alpha_1(t), \alpha_2(t)$), with period T , that results in a net swimming ($x(t), y(t)$) or turning $\beta(t)$ of the fish body.

This model is a canonical example of a class of under-actuated control problems whose dynamics can be described over the shape space using tools from geometric mechanics [44,55–57]. On the right-hand side of Eq. (4), the first term represents a dynamic phase or drift and the second term represents a geometric phase over the fish shape space (α_1, α_2) [45]. The geometric phase is best described in terms of the local connection matrix \mathbb{A} [45,54], which is a function only of the shape variables α_1 and α_2 ,

$$\mathbb{A} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \\ A_{\beta 1} & A_{\beta 2} \end{bmatrix} := -\mathbb{I}_{\text{lock}}^{-1} \mathbb{I}_{\text{couple}}. \quad (5)$$

Each row of \mathbb{A} describes a nonlinear vector field over the shape space, giving rise to three vector fields $\mathbf{A}_1 \equiv (A_{11}, A_{12})$, $\mathbf{A}_2 \equiv (A_{21}, A_{22})$, and $\mathbf{A}_\beta \equiv (A_{\beta 1}, A_{\beta 2})$ over the (α_1, α_2) plane as shown in Fig. 2(a). In driftless systems, net locomotion is fully controlled by the fish shape changes as dictated by the connection matrix \mathbb{A} . However, in the presence of drift, shape control is not sufficient to determine the fish motion in the physical space, which is then affected by the drift term in Eq. (4).

III. GEOMETRIC PHASES

Geometric phases are defined as the net locomotion (x, y, β) that results from prescribed cyclic shape changes in the (α_1, α_2) plane at zero total momentum (no drift). Inverting Eqs. (1) and substituting Eq. (5) into Eq. (4) at zero total momentum, we arrive at

$$\begin{bmatrix} \cos \beta & \sin \beta & 0 \\ -\sin \beta & \cos \beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\beta} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \\ A_{\beta 1} & A_{\beta 2} \end{bmatrix} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (6)$$

Motions (x, y, β) in the physical space are obtained by integrating Eq. (6) with respect to time. Rotations are directly proportional to the line integral of the vector field \mathbf{A}_β as evident by integrating the last equation in Eq. (6),

$$\beta(T) - \beta(0) = \oint_C d\beta = \oint_C (A_{\beta 1} d\alpha_1 + A_{\beta 2} d\alpha_2). \quad (7)$$

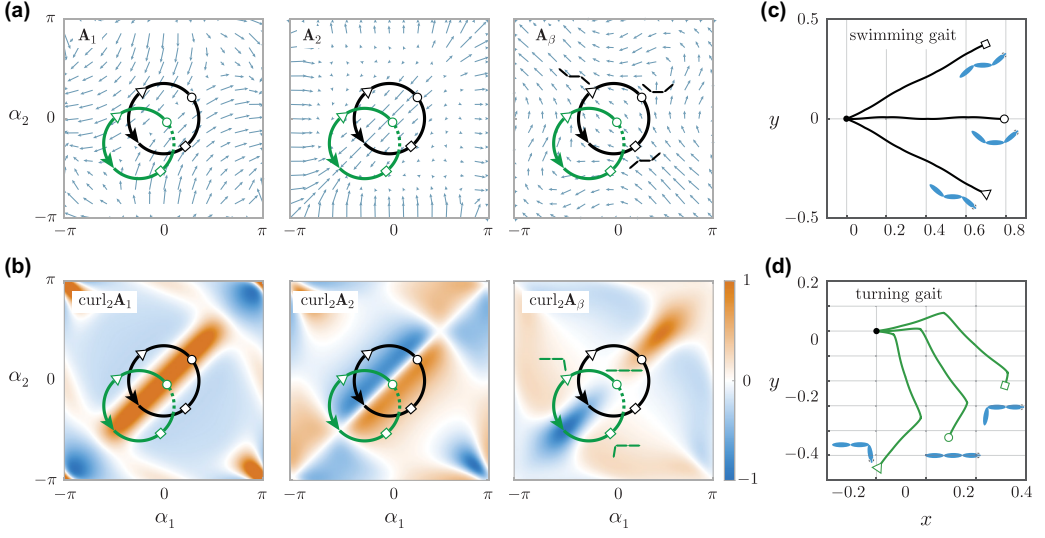


FIG. 2. Using connection matrix for simple gait design. (a), (b) Rows of connection matrix \mathbf{A} give us three vector fields, $\mathbf{A}_1 \equiv \partial v_1 / \partial \alpha_i$, $\mathbf{A}_2 \equiv \partial v_2 / \partial \alpha_i$, and $\mathbf{A}_\beta \equiv \partial \beta / \partial \alpha_i$, the curl of which yields three corresponding scalar fields. Magnitude of the scalar fields are normalized to be within $[-1, 1]$. Value of the scalar fields can facilitate the design of simple swimming and turning gaits, as shown by black and green circles, respectively. Six body configurations, each corresponding to points marked by black and green \circ , \triangle , and \square , are sketched for additional clarity. (c), (d) Fish that start with body centered at the origin of the x - y plane and follow the same gait circle swim and/or turn in different directions in the physical space depending on initial body shapes; the initial shape is depicted in blue at the end of each trajectory.

Here, T is the time period for going around the closed trajectory in the shape space once. Using Green's theorem, we get

$$\beta(T) - \beta(0) = \iint \left(\frac{\partial A_{\beta 2}}{\partial \alpha_1} - \frac{\partial A_{\beta 1}}{\partial \alpha_2} \right) d\alpha_1 d\alpha_2 = \iint \text{curl}_2(\mathbf{A}_\beta) d\alpha_1 d\alpha_2. \quad (8)$$

Here, curl_2 denotes the two-dimensional curl as a scalar field over the (α_1, α_2) plane. The scalar field $\text{curl}_2(\mathbf{A}_\beta)$ provides an intuitive tool for understanding the effect of a cyclic shape deformation on the net rotation of the fish: net rotations are proportional to the integral of $\text{curl}_2(\mathbf{A}_\beta)$ over the area enclosed by a closed shape trajectory (see Fig. 2). However, translational motions (x, y) are not directly proportional to the area integrals of $\text{curl}_2(\mathbf{A}_1)$ and $\text{curl}_2(\mathbf{A}_2)$, but to a combination of all three integrals coupled through the fish rotational dynamics as evident from Eq. (6). This coupling is clear when the equations of motion for x and y in Eq. (6) are written in scalar form,

$$\begin{aligned} \dot{x} &= [A_{11}\dot{\alpha}_1 + A_{12}\dot{\alpha}_2] \cos \beta - [A_{21}\dot{\alpha}_1 + A_{22}\dot{\alpha}_2] \sin \beta, \\ \dot{y} &= [A_{11}\dot{\alpha}_1 + A_{12}\dot{\alpha}_2] \sin \beta + [A_{21}\dot{\alpha}_1 + A_{22}\dot{\alpha}_2] \cos \beta. \end{aligned} \quad (9)$$

Despite this complication, the scalar fields defined by $\text{curl}_2(\mathbf{A}_1)$, $\text{curl}_2(\mathbf{A}_2)$, and $\text{curl}_2(\mathbf{A}_\beta)$, shown in Fig. 2(b), are informative of the net translational (x, y) and rotational β motions of the fish. To illustrate the utility of these curl fields, we show two examples of cyclic shape changes depicted in black and green lines. A fish changing its shape following the black line undergoes zero net rotation because the area integral of $\text{curl}_2(\mathbf{A}_\beta)$ is identically zero, but it swims forward in the (x, y) plane. The net displacement per period is a conserved quantity, whereas the direction of motion depends on a combination of the fish initial shape $(\alpha_1(0), \alpha_2(0))$ and initial orientation $\beta(0)$ as shown in Fig. 2(c). Here, we consider three different initial shapes, depicted by the markers \circ , \triangle , and \square ,

all at $\beta(0) = 0$. Similarly, shape deformations following the green line lead to net reorientations in the physical space, as shown in Fig. 2(d). Evidently, no net motion occurs if the shape trajectory is degenerate, that is, does not enclose an area in the shape space. Furthermore, a rescaling of time does not affect the net motion of the fish, only the speed at which the fish completes these cyclic shape changes.

In Fig. 2 and hereafter, the equations of motion are nondimensionalized using the total length of the fish and the total mass in the head-to-tail direction of the straight fish as the characteristic length and mass scale. Specifically, we set the total mass of the fish body to be equal to the added mass (actual mass of the fish is negligible). We leave the time scale unchanged because the characteristic time depends on the speed of shape changes, which is a control parameter to be determined by the controller.

The scalar fields $\text{curl}_2(\mathbf{A}_1)$, $\text{curl}_2(\mathbf{A}_2)$, and $\text{curl}_2(\mathbf{A}_\beta)$ over the shape space encode information about the net locomotion of the fish in a driftless environment, and can be used to design simple swimming and turning gaits as shown in Fig. 2. However, these geometric tools do not allow for a straightforward design of control policies for arbitrary motion planning [52,54], and they are even less instructive in the presence of drift. Next we consider an RL driven approach for motion control.

IV. MOTION CONTROL VIA REINFORCEMENT LEARNING

We use RL to train the three-link fish on two different tasks: (i) to swim parallel to a desired direction in a driftless environment and (ii) to swim towards a target point located at a distance ρ and angular position ψ from the fish nose, with ρ and ψ expressed in the fish frame of reference [Fig. 1(b)]. Given the rotational symmetry of the fish-fluid space, in the first task, we fix the desired direction to be parallel to the x -axis without loss of generality. For the second task, we first train the fish in a driftless environment, then introduce drift and train again in the presence of drift. The first task allows for direct comparison of the performance of the trained policy to manually designed policies in the context of geometric mechanics as described in Sec. III. The second task allows for evaluation and comparison of the performance of the driftless and drift-aware policies under environmental perturbations.

Central to any RL implementation are the notions of the state of the system, the observations given to the learning agent, the actions taken by the agent, and the rewards given to the agent in light of its behavior. The state s_t of the fish-fluid-target system at a time t is given by the fish position and orientation in inertial frame (x, y, β) , its shape (α_1, α_2) , and the target position relative to the fish, (ρ, ψ) [see Figs. 1(b) and 1(c)]. As sensory input, we provide the fish a set of observations based on its proprioception of its own shape α_1 and α_2 , as well as an egocentric observation of the task; namely, for controlling the direction of swimming, the fish knows the desired swimming direction relative to itself $\psi = -\beta$, and for swimming towards a target, it knows the relative angular position ψ of the target point. This yields a set of observations $o_t = (\alpha_1, \alpha_2, \psi)_t$. Additionally, when training in the presence of drift, the magnitude and direction of the drift vector (p_x, p_y) are also provided as observations. As control action, the fish has direct control of its shape using the rate of shape changes as actions $a_t = (\dot{\alpha}_1, \dot{\alpha}_2)_t$. With this choice of action, the control can be projected onto the shape space and directly compared to the geometric mechanics approach. We constrain the value of the actions to be between -1 and 1 rad per unit time, and we impose limits on the joint angle so that α_1 and α_2 are only allowed to change between $-2\pi/3$ and $2\pi/3$ rad to avoid self-intersection.

In RL, the decision making process is modeled as a stochastic control policy $\pi_\theta(a_t|o_t)$ that produces actions a_t given observations o_t of the state of the fish-fluid system. The policy is parametrized by a set of parameters θ to be optimized. An optimal policy is learned to produce behavior that maximizes rewards. We use a dense shaping reward; that is, the fish is given a reward at every decision time step. Specifically, we set the reward to be the distance the fish travels in the desired direction towards the target state. For learning to swim parallel to the x -axis, we use the reward $r_t = x_{\text{nose},t+1} - x_{\text{nose},t}$, which is the change in the fish nose position along the x -axis. Note that this choice of reward function breaks the head-tail symmetry of the fish favoring higher

reward when turning towards the positive x -axis. For learning to swim towards a target, the reward $r_t = \rho_t - \rho_{t+1}$ is based on the change in the relative distance ρ from the fish nose to the target. The return $R_t = \sum_{t'}^{\infty} \gamma^{t'-t} r_{t'}$ is defined as the infinite horizon objective based on the sum of discounted future rewards, where $\gamma \in [0, 1]$ is known as the discount factor; it determines the preference for immediate over future rewards. We set $\gamma = 0.99$ to make the fish foresighted. The goal is to arrive at an optimal set of parameters θ that maximizes the *expected* return $\mathcal{J}(\pi_\theta) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_t]$ for a distribution of initial states. Here, the expectation is taken with respect to the distribution over trajectories $\pi(\tau)$ induced jointly by the fish dynamics, viewed as a partially observable Markov decision process, and the policy $\pi_\theta(a_t|o_t)$. One approach to solving this optimization problem is to use a policy gradient method that computes an estimate of the gradient $\nabla_\theta \mathcal{J}$ for learning. Policy gradient methods are widely used to learn complex control tasks and are often regarded as the most effective reinforcement learning techniques, especially for robotics applications [60–64]. Here, we use a specific class of policy gradient methods, known as actor-critic methods [65,66] where the fish learns simultaneously a policy (actor) and a value function (critic). We implement this method using the clipped advantage proximal policy optimization (PPO) algorithm proposed in [67]. This algorithm ensures fast learning and robust performance by limiting the amount of change allowed for the policy within one update. A pseudocode implementation of the PPO algorithm and additional implementation details are provided in Appendix B.

V. TRAINING THE FISH TO SWIM

We trained the three-link fish to (i) swim parallel to the x -direction in a driftless environment and (ii) swim towards a target point in the absence and presence of drift. We refer to the first task as *direction control* for short, and the second task as *naive* and *drift-aware target seeking*, respectively, based on their awareness of drift.

For the purpose of efficient training we imposed a finite time interval, following which the system state was reset to the initial state for a new round of training. Each round is referred to as an *episode*. In all training episodes, we initialized the fish center to be at the origin of the inertial frame, and we initialized the shape angles α_1, α_2 and body orientation β by sampling from a uniform distribution over all permissible angles to maximize the chances for robust learning. We fixed the maximum episode length to 150 time steps, with no early termination allowed. In training the target-seeking policies, the target was initially placed at a fixed distance (three-unit length) from the fish center but at a random orientation. For the drift-aware policy, a drift was introduced in the form of a nonzero total linear momenta p_x, p_y in Eqs. (4). The drift magnitude was sampled from a uniform distribution between zero and 0.15 and its direction from a uniform distribution between zero and 2π , such that the drift was kept constant within a training episode. Based on the nondimensional scales introduced in Sec. III, one unit of drift aligned parallel to a straight fish causes it to translate one unit length per unit time.

For the direction-control policy, we performed 24 runs of RL training with 10 000 episodes in each run. The training process is illustrated in terms of reward evolution, calculated by taking the sum of all rewards in an episode in Appendix B, Fig. 8(a). There was some variability among the seeds, but most trained policies performed very well; only a single policy did not converge by the end of training episodes. Note that fluctuations in reward after policy convergence are partly due to the stochasticity built into the policy itself and partly due to variation in task difficulty given the random initial conditions: different initial conditions require different amounts of time and effort for the fish body to align with the x -axis.

It is worth pointing out here the training results of the direction-control policy were affected significantly by the episode length. In order to swim in a desired direction starting from an arbitrary initial orientation, the fish has to first turn in that direction, then swim forward. Policies trained with longer episodes performed better in the swimming portion of the trajectory but failed to make large-angle turns, as training data collected on swimming significantly outweighed those collected on turning. On the other hand, policies trained with shorter episodes made turns of any angle, but

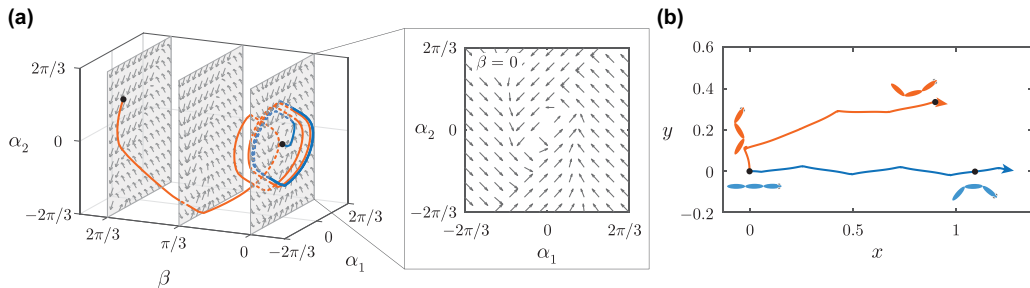


FIG. 3. Visualizing the direction-control RL policy. (a) Given the direction-control task, we visualize the mean RL policy actions $(\dot{\alpha}_1, \dot{\alpha}_2)$ as vector fields in the observation space of $(\alpha_1, \alpha_2, \beta)$. Two example observation trajectories starting at $\beta(0) = 0, \alpha_1(0) = 0, \alpha_2(0) = 0$ (blue) and $\beta(0) = 2\pi/3, \alpha_1(0) = -\pi/3, \alpha_2(0) = \pi/3$ (orange) are plotted with slices of the mean policy at $\beta = 0, \pi/3, 2\pi/3$. The inset shows a flattened view of the slice at $\beta = 0$. (b) Physical space trajectories of the corresponding examples shown in (a) are plotted together with the fish body configurations at the starting point and chosen points near the ends.

were less likely to swim straight after turning. We chose the episode length, following several trial-and-error trainings, to be 150 as a reasonable compromise to learn to both turn and swim effectively.

The evolution of rewards during training of the two target-seeking policies are plotted in Fig. 8(b), each with 20 runs of the RL algorithm and 15 000 episodes in each run. The naive policy converged faster than the drift-aware policy, and both policies converged slower than the direction-control policy. These results indicate that the task itself, as well as variations in the environment and number of observations, affects the convergence rate, that is, the learning difficulty. Note that the numerical value of the reward is not directly comparable between policies for different tasks.

We evaluated the performance of the trained policies by testing them under two types of conditions: conditions similar to those seen during training and perturbed conditions not seen during training, as discussed next.

VI. BEHAVIOR OF TRAINED FISH

To visualize the RL direction-control policy, we plot in Fig. 3(a) the action vector fields $(\dot{\alpha}_1, \dot{\alpha}_2)$ over the observation space $(\alpha_1, \alpha_2, \beta)$. These vector fields depend on the orientation β of the fish in the physical space such that the control policy $(\dot{\alpha}_1, \dot{\alpha}_2)$ forms a *foliation* over β . Three slices of this foliation are highlighted. The right-hand side of Fig. 3(a) provides a closer look at the policy slice at $\beta = 0$; the arrows indicate the mean actions advised by the policy for a given set of observations α_1, α_2 at $\beta = 0$. In Fig. 3(b), we show the details of two trajectories in the physical space starting from two distinct configurations. In the first test, the fish starts at zero orientation, $\beta(0) = 0$, in a straight shape, $(\alpha_1(0), \alpha_2(0)) = (0, 0)$. The goal of the fish is simply to swim forward. In the second test, the initial orientation and shape are $\beta(0) = 2\pi/3$ and $(\alpha_1(0), \alpha_2(0)) = (-\pi/3, \pi/3)$, from which the fish needed to turn and swim along the x -axis. In both cases, the fish is able to turn to the desired direction and swim steadily. In Fig. 3(a), we highlight the corresponding trajectories in the $(\alpha_1, \alpha_2, \beta)$ space. As the fish moves through the physical space, β changes, causing the fish to take actions from distinct slices of the foliation of action vector fields. Both trajectories tend to the same periodic swimming cycle around $\beta = 0$, indicating that the control actions are similar once both fish are aligned with the desired swimming direction.

We further explore the shape changes undertaken by the second (red) fish by superimposing these shape changes onto the scalar fields $\text{curl}_2 \mathbf{A}_1$, $\text{curl}_2 \mathbf{A}_2$ and $\text{curl}_2 \mathbf{A}_\beta$ introduced in Sec. III [see Fig. 4(a)]. The corresponding motion in the physical space is depicted in red in Fig. 4(b). The shape deformations produced by the RL policy can be interpreted as consisting of two regimes: an initial




FIG. 4. RL provides smooth transition between turning and swimming gaits. (a) The shape space trajectory of the fish reorienting itself to swim parallel to the x -axis produced by mean RL policy is superimposed to the scalar curl fields from Fig. 2(b). Note that this trajectory starts off-centered and smoothly moves to cycles that are symmetric about the origin. (b) The physical space trajectory due to the mean RL policy (red) in comparison to a manually patched turning-to-swimming trajectory (green and black) using circular gaits in Fig. 2. Without further fine-tuning on the shape of the gaits, the manually patched result shows a more abrupt and unnecessary turn angle. In both simulations, the fish start in a straight configuration with an orientation of $\beta(0) = \pi/3$ at the origin.

turning regime followed by a forward swimming regime. The turning regime is indicated by the initial portion of the shape trajectory enclosing most of the blue area in the $\text{curl}_2\mathbf{A}_\beta$ image; the integral in Eq. (7) along this portion of the trajectory is negative, leading to a clockwise rotation. The swimming regime is indicated by the periodic shape changes enclosing the rectangular orange portion in the $\text{curl}_2\mathbf{A}_1$ image. The area integral of $\text{curl}_2\mathbf{A}_1$ along this portion of the trajectory is positive, whereas the corresponding area integrals of $\text{curl}_2\mathbf{A}_2$ and $\text{curl}_2\mathbf{A}_\beta$ are identically zero, leading to net motion in the positive x -direction. These shape deformations and resulting turning and swimming motions can be compared to a manually designed shape trajectory based on the turning (green) and swimming (black) gaits in Figs. 2(a) and 2(b). Specifically, starting from a straight fish configuration, we follow the solid portion of the green trajectory (turning) in Figs. 2(a) and 2(b), and transition into the black trajectory (swimming) at the second intersection of the green and black shape trajectories. The resulting motion in the physical space is superimposed onto Fig. 4(b). Both the RL and manually patched gaits lead the fish to turn and swim parallel to the x -axis; however, in the RL produced motion, the transition between turning and swimming is smoother. Note that, in the swimming regime, the RL policy produces cyclic shape deformations that do not maximize the area integral of $\text{curl}_2\mathbf{A}_1$ (shape trajectory does not enclose the whole orange portion in the $\text{curl}_2\mathbf{A}_1$ image). Indeed, maximizing this integral is not optimal for forward swimming as discussed next. The key lies in the fact that the physics of the problem, specifically, the rotational motion of the fish, couples displacements in the x - and y -directions. Our model-free RL policy captures this effect with no explicit knowledge of the physics of the system.

To better explain this optimality of the RL policy, we manually prescribed cyclic shape changes that follow rectangular trajectories reminiscent of the trajectory generated by the RL policy for forward swimming. The manually designed shape trajectories all share the same width as the RL policy albeit at different lengths, namely, 1.2 (small), 2.4 (medium), and 3.6 (large) to enclose increasingly larger regions of the orange portion of the $\text{curl}_2\mathbf{A}_1$ (see right column of Fig. 5). These cyclic deformations result in net displacements in the x -direction with zero-mean excursions in the y -direction. The y -excursions are due to the fact that, even though the area integrals of $\text{curl}_2\mathbf{A}_2$ and $\text{curl}_2\mathbf{A}_\beta$ over the regions enclosed by these shape trajectories are identically zero, leading to zero net rotation over a full cycle of shape deformations, the instantaneous rotations β of the fish body couple displacements in the x - and y -directions, as evident from Eqs. (9). For the cyclic shape changes in Fig. 5(a), the amplitude of y -excursions is small but so is the net displacement in the x -direction; meanwhile in Fig. 5(c), both the net displacement in the x -direction and the amplitude of the

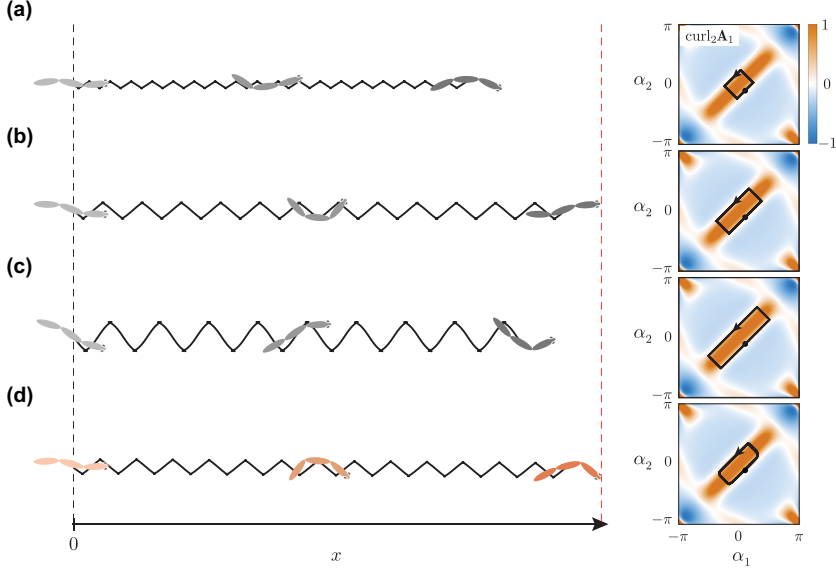


FIG. 5. Racing against the RL fish. To test the optimality of our direction-control RL policy, we compare forward swimming performance between the geometrically designed gaits and RL results. In (a)–(c) we use rectangular gaits with length equal to 1.2 (small), 2.4 (medium), and 3.6 (large), respectively, and width all equivalent to shape space trajectory following the mean actions from RL policy in (d). All shape space trajectories are shown superimposed on top of the $\text{curl}_2\mathbf{A}_1$ field on the right. Physical space trajectories on the left show that the mean RL policy achieves its excellent performance by choosing an optimal amount of lateral oscillation during forward swimming, while the small and large rectangular gaits move slower due to either insufficient or overwhelming sideways motion. Note that fish in (a)–(c) are initialized with the same shape but slightly different initial orientations to ensure they all swim in exactly the x -direction. In addition, all fish utilize the maximum actions allowed at each time step.

y -excursions are large. The fastest fish is the one that maximizes forward motion while minimizing lateral movements, as shown in Fig. 5(b) and recapitulated in the RL result shown in Fig. 5(d). It is worth emphasizing that the RL policy arrives at this optimal solution merely by sampling observations, actions, and rewards, with no prior or developed knowledge of the physics of the problem.

Next, we investigated the effect of the initial orientation $\beta(0) \in [-\pi, \pi]$ on the amount of control effort required to turn and swim parallel to the x -axis. Figure 6(a) shows three examples of fish following the trained policy starting from three distinct initial orientations $\beta(0) = 0, \pi/3$, and $2\pi/3$ and a straight configuration centered at the origin of the physical space. To measure the actuation effort needed in these motions, we used the integral $\int_0^t T_{\text{shape}} dt$ of the actuation energy $T_{\text{shape}} = \frac{1}{2}(J + m_2 a^2)(\dot{\alpha}_1^2 + \dot{\alpha}_2^2)$, which is the energy imparted to the fluid by the fish shape changes (see Appendix A). Figure 6(b) shows the actuation effort versus the initial orientation of a straight fish. Here the fish was instructed by the stochastic policy with the same action noise as in the training process. The actuation effort, as well as its variation due to noise, is larger for larger turning angles.

Lastly, we examined the behavior and effort of a fish swimming instructed to reach a known target in an environment with zero drift. Figure 6(c) shows examples of the fish swimming motion in the physical space for targets located at $\psi = \pi/6$ and $\psi = 5\pi/6$. All tests ran until the fish reached the target or a maximum interval of 1000 time steps was exceeded. We varied the target angular position while maintaining the fish initial shape and orientation (the fish always started straight in the x -direction), and we calculated the actuation effort as a function of the target orientation [see Fig. 6(d)]. The actuation effort increases as the target moves from the front to the back of the fish,

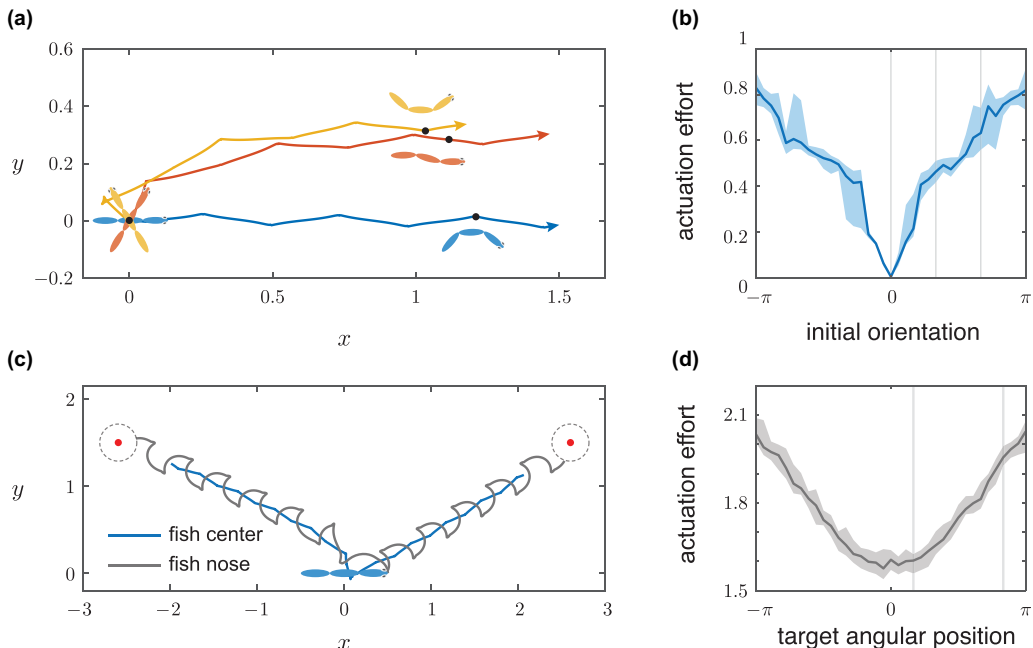


FIG. 6. Performance of RL policies in a driftless environment. (a) Swimming trajectories (center of mass) arrived at by the mean action of the direction-control policy are shown for fish starting at orientations $\beta = 0$ (blue), $\pi/3$ (red), and $2\pi/3$ (yellow). (b) The actuation effort of reorientation is roughly proportional to the absolute value of the initial fish orientation due to the amount of turning maneuvers required. Results shown are based on 25 stochastic policy rollouts per tested orientation angle. (c) Center of mass (blue) and nose (grey) swimming trajectories using the mean action of the naive target-seeking policy are shown with two targets located at an angular position of $\pi/6$ and $5\pi/6$. The fish is considered to have reached the target when its nose is within $\epsilon = 0.2$ distance from the target (dotted circles). (d) Actuation effort needed to perform shape increases as the target angular position changes away from zero. Results shown use 25 stochastic policy rollouts per tested target angle. Note that solid lines and shaded regions of (b) and (d) show the median results and variations between 25th and 75th percentile, respectively. In (b), performance is calculated based on the actuation effort it takes for the fish to turn in the x -direction, while in (d) it is based on the effort it takes to reach the target.

because it requires larger turns in order for the fish to align its heading direction with the direction of the target; this is consistent with our findings based on the direction-control policy. It is worth emphasizing that the direction-control task is equivalent to the target-seeking task with the target placed at $x = +\infty$. It is also worth noting that the reward function is measured from the nose of the fish, thus breaking the fore-aft symmetry of the fish, and leading to control action that favors turning and swimming headfirst towards the target.

We tested the behavior and effort of the target-seeking fish in the presence of nonzero drift. Figures 7(a) and 7(c) show a comparison between the naive and drift-aware policies for targets located at $\psi = \pi/6$ and $\psi = 5\pi/6$, with drift of magnitude of 0.1 pointing to the x -direction and the $-y$ -direction, respectively. The naive policy (grey lines) is able to reach the target, even though it does not directly observe the drift, albeit following different actions and swimming trajectories than the drift-aware policy (orange lines). Specifically, when following the drift-aware policy, the fish curls less when the drift is helpful and curls more when the drift is unfavorable. We assessed the performance of the two policies for various drift magnitudes and directions. In Fig. 7(b), we calculated the actuation effort as a function of the drift magnitude with a fixed drift direction for two target locations. The naive policy outperforms the drift-aware policy for small drift (drift magnitudes

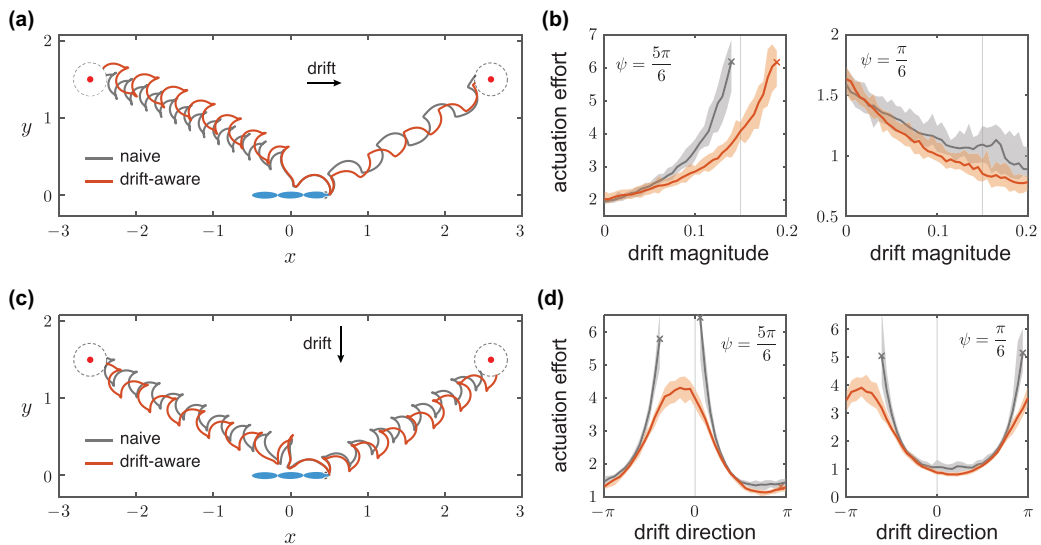


FIG. 7. Performance of target-seeking policies in the presence of drift. We compare the swimming trajectories and actuation efforts of target-seeking policies for different drift magnitudes and direction. (a), (c) Naive policy (grey) and drift-aware policy (orange) produce similar average actions in environments with constant (0.1) drift in the positive x -direction and negative y -direction, respectively. Both panels showcase trajectories reaching targets located three unit lengths away with angular position ψ of $\pi/6$ and $5\pi/6$. (b) For a fixed drift direction (positive x -direction), actuation effort as a function of drift magnitude evolves differently depending on target angular position ψ . When drift is against the direction of the target (left), both policies fail to reach the target on average for large enough drift (cross marks). No failures are observed when drift facilitates swimming towards the target (right). In both cases, the drift-aware policy significantly outperforms the naive policy at large drift magnitudes by saving actuation effort. Intriguingly, inclusion of extra observations in drift-aware policies seems to result in slightly suboptimal performance when the drift magnitude is very small. (d) With the drift magnitude fixed to 0.15, the naive policy fails to reach the target if the drift direction is near the opposite end of the target angular position ψ . However, at this drift magnitude, the drift-aware policy can still reach both targets regardless of the drift direction. Note that solid lines and shaded regions of (b) and (d) show the median results and variations between 25th and 75th percentile, respectively.

less than 0.025) even when the drift is adversarial, but the naive policy loses or even fails to finish the task when the drift is large, especially when the drift is in the adverse direction to the target location. This implies that it might be wise to discard some sensory input (observations) when the perturbation due to drift is weak, especially if these extra observations act more like a distraction than a clue. But as the perturbation gets stronger, it is necessary to take more observations into account. Both the naive and drift-aware policies are not able to complete the task in the given amount of time when the drift magnitude is very large and its direction is adversarial to the target location. This is because the shape actuation has no direct control over the drift itself. In Fig. 7(d), we fixed the magnitude of drift to 0.15 and changed its direction. Using the actuation effort as our performance metric as before, the drift-aware policy has better or similar performance under all tested conditions. The naive policy fails when the drift acts in the adverse direction relative to the target while the drift-aware policy is always able to reach the target before the episode terminates.

VII. CONCLUDING REMARKS

We considered a three-link fish swimming in potential flow. We reviewed that swimming in potential flow can be expressed as a combination of a dynamic phase (drift) and geometric phase

(driftless) over the shape of fish body deformation [44]. In the driftless case, net locomotion is purely determined by the fish shape deformations, and geometric techniques can be used for gait design and motion planning over the fish shape space [52,54], but shape actuation cannot control the drift itself. Yet, even in the driftless regime, motion planning starting from arbitrary fish orientation and shape is not trivial. In this paper, we applied model-free reinforcement learning techniques for controlling the fish motion, and we arrived at optimal policies for swimming (i) in a desired direction and (ii) towards a target in the absence and presence of drift. The RL-based policies produce behavior that is robust to variations in the fish initial shape and orientation and target location. We used the actuation effort as a measure of the policy performance under various initial conditions and in various environments, and we quantified the robustness of the RL policies to the presence of drift. We found that although the fish has no control over the drift itself, the fish learns to take advantage of the presence of moderate drift to reach its target. Therefore, it might be wise to discard some sensory input (observations) for weak signals because these extra observations could act more like a distraction than a clue. We also found that for small drift, the drift-naive policy outperforms the drift-aware policy even when drift is adversarial to the target location. However, large adversarial drift hinders the fish ability to locate the target. Importantly, these insights into the RL policies were achievable by combining tools from geometric mechanics with RL-based control. Geometric tools such as the concept of shape space, combined with the RL notion of the action space, provide a useful and novel framework for visualizing and interpreting the RL policies as action vector fields over the shape space.

A few comments on the advantages and limitations of our implementation are in order. Despite algorithmic advantages, obtaining an RL policy is computationally costly, especially when the environment simulator involves high-fidelity fluid-structure interaction models. To circumvent this problem, recent work on training fish to swim uses a limited set of observations and actions [39]. For example, the zebrafish model of [39] allows only five discrete actions, each corresponding to a fixed amplitude of body curvature change. Reduced-order fluid models, such as the potential flow model used here, offer an enticing framework for designing control laws that can later be tested and refined using more realistic flow environments, as done in [20] for a manually designed swimming gait in [44]. Specifically, in the simplified potential flow environment employed here, we are able to train continuous action policies using a rich set of observations, with an eye on probing the performance of these policies in more realistic flow environments in future work.

Lastly, we invite interested readers to visit our source repository for the complete details of our implementation [68].

ACKNOWLEDGMENTS

E.K. would like to acknowledge support from the Office of Naval Research through Grants No. N00014-17-1-2287 and No. N00014-17-1-2062. The authors are thankful for useful discussions with Dr. Yuval Tassa.

APPENDIX A: PHYSICS OF THE FISH MODEL

We review the derivation of the equations of motion governing the swimming of an articulated three-link fish in potential flow (Fig. 1).

1. Fish kinematics

Consider planar motions of the three-link fish. Let $\mathbf{x} = (x, y)$ denote the position of the center of mass G of the middle link, and let β denote the orientation of the fish relative to a fixed inertial frame, here taken to be the angle between the x -axis and the major axis of symmetry of the middle link. Let α_1 and α_2 be the rotation angles of the front link relative to the middle link and the middle link relative to the rear link; that is to say, (α_1, α_2) represents the shape of the three-link fish. It is convenient for the following development to introduce a body-fixed frame $(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$, attached

at G and co-rotating with the middle link. This body-fixed frame is related to the inertial frame $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ via a rigid-body rotation such that $\mathbf{e}_1 = \cos \beta \mathbf{b}_1 - \sin \beta \mathbf{b}_2$, $\mathbf{e}_2 = \sin \beta \mathbf{b}_1 + \cos \beta \mathbf{b}_2$, and $\mathbf{e}_3 = \mathbf{b}_3$.

The velocity (\dot{x}, \dot{y}) of the center of mass of the middle link, when expressed in the body-fixed frame, is given by

$$\mathbf{v} = v_1 \mathbf{b}_1 + v_2 \mathbf{b}_2 = (\dot{x} \cos \beta + \dot{y} \sin \beta) \mathbf{b}_1 + (-\dot{x} \sin \beta + \dot{y} \cos \beta) \mathbf{b}_2. \quad (\text{A1})$$

Assuming all three links are made of identical ellipsoids of half length a , half width b , and half height c , the velocities of the centers of mass G_1 and G_2 of the front and rear link, expressed in the body-fixed frame of the middle link, are given by ($i = 1, 2$, denote the front and rear link, respectively)

$$\mathbf{v}_i = (v_1 \mp a\dot{\alpha}_i \sin \alpha_i - a\dot{\beta} \sin \alpha_i) \mathbf{b}_1 + (v_2 \pm a\dot{\beta} + a\dot{\alpha}_i \cos \alpha_i \pm a\dot{\beta} \cos \alpha_i) \mathbf{b}_2. \quad (\text{A2})$$

The angular velocities of the middle, front, and rear links are given by $\dot{\beta}$, $\dot{\beta} + \dot{\alpha}_1$, and $\dot{\beta} - \dot{\alpha}_2$, respectively. For our simulations, we used $a:b:c = 5:1:5$ as the geometry of the ellipsoids.

2. Kinetic energy of the articulated body

In the absence of the fluid, the kinetic energy of the articulated three-link body is given by

$$T_{\text{body}} = \frac{1}{2} m_s \mathbf{v} \cdot \mathbf{v} + \frac{1}{2} J_s \dot{\beta}^2 + \frac{1}{2} \sum_i [m_s \mathbf{v}_i \cdot \mathbf{v}_i + J_s (\dot{\beta} \pm \dot{\alpha}_i)^2], \quad (\text{A3})$$

where $m_s = \frac{4}{3} \pi abc \rho_s$ and $J_s = \frac{1}{5} (a^2 + b^2) m_s$ are the mass and moment of inertia of each solid link with ρ_s the density of the links.

3. Kinetic energy of the fluid

The three-link fish is submerged in an unbounded domain of incompressible and irrotational fluid, such that the fluid velocity $\mathbf{u} = \nabla \phi$ can be expressed as the gradient of a potential function ϕ . It is a standard result in potential flow theory that the kinetic energy of the fluid can be expressed in terms of the variables of the submerged solid [44,45,59]. In the case of a single ellipsoid, the kinetic energy of the fluid is given by $T_{\text{fluid}} = [(m_{1a} v_1^2 + m_{2a} v_2^2) + J_a \dot{\beta}^2]/2$, where m_{1a} , m_{2a} , and J_a are the added mass and added moment of inertia due to the presence of the fluid, expressed in a body-fixed frame that coincides with the major and minor axes of the ellipsoid. These quantities depend on the geometric properties a, b, c of the submerged ellipsoid, as are given in Appendix B of [69]. For a nonspherical body, the added masses m_{1a} , m_{2a} depend on the direction of motion: the added mass is larger when moving in the direction of the minor axis of symmetry of the ellipsoid, that is to say, in the transverse direction; hence $m_{1a} \leq m_{2a}$.

In the case of the three-link fish, the kinetic energy of the fluid is of the form

$$\begin{aligned} T_{\text{fluid}} &= \frac{1}{2} m_{1a} v_1^2 + \frac{1}{2} m_{2a} v_2^2 + \frac{1}{2} J_a \dot{\beta}^2 + \frac{1}{2} \sum_i J_a (\dot{\beta} \pm \dot{\alpha}_i)^2 \\ &+ \frac{1}{2} \sum_i m_{1a} (v_1 \cos \alpha_i \pm v_2 \sin \alpha_i + a\dot{\beta} \sin \alpha_i)^2 \\ &+ \frac{1}{2} \sum_i m_{2a} (\mp v_1 \sin \alpha_i + v_2 \cos \alpha_i \pm a\dot{\beta} \cos \alpha_i \pm a\dot{\beta} + a\dot{\alpha}_i)^2. \end{aligned} \quad (\text{A4})$$

Here we transform the velocity components of head and tail by α_1 and α_2 , respectively, to match the added mass components.

4. Kinetic energy of the body-fluid system

The kinetic energy of the fish-fluid system is obtained by taking the sum of Eqs. (A3) and (A4), which can be expressed in matrix form as follows:

$$T = T_{\text{body}} + T_{\text{fluid}} = \frac{1}{2} \begin{bmatrix} v_1 \\ v_2 \\ \dot{\beta} \\ \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}^T \left[\begin{array}{c|c} \mathbb{I}_{\text{lock}} & \mathbb{I}_{\text{couple}} \\ \hline \mathbb{I}_{\text{couple}}^T & \mathbb{I}_{\text{shape}} \end{array} \right] \begin{bmatrix} v_1 \\ v_2 \\ \dot{\beta} \\ \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (\text{A5})$$

Here, \mathbb{I}_{lock} is a 3×3 locked mass matrix, function of α_1 and α_2 ,

$$\mathbb{I}_{\text{lock}} = \begin{bmatrix} \mathbb{M} & \mathbb{H} \\ \mathbb{H}^T & \mathbb{J} \end{bmatrix}, \quad (\text{A6})$$

where \mathbb{M} is a 2×2 mass matrix given by

$$\mathbb{M} = \begin{bmatrix} m_1(1 + \sum_i \cos^2 \alpha_i) + m_2 \sum_i \sin^2 \alpha_i & \frac{1}{2}(m_1 - m_2)(\sin 2\alpha_1 - \sin 2\alpha_2) \\ \frac{1}{2}(m_1 - m_2)(\sin 2\alpha_1 - \sin 2\alpha_2) & m_2(1 + \sum_i \cos^2 \alpha_i) + m_1 \sum_i \sin^2 \alpha_i \end{bmatrix}, \quad (\text{A7})$$

\mathbb{J} is a moment-of-inertia scalar given by

$$\mathbb{J} = 3J + m_1 a^2 \sum_i \sin^2 \alpha_i + m_2 a^2 \sum_i (1 + \cos \alpha_i)^2, \quad (\text{A8})$$

and \mathbb{H} is given by

$$\mathbb{H} = \begin{bmatrix} \frac{1}{2}(m_1 - m_2)a \sum_i \sin 2\alpha_i - m_2 a \sum_i \sin \alpha_i \\ \frac{1}{2}(m_1 - m_2)a(\cos 2\alpha_2 - \cos 2\alpha_1) + m_2 a(\cos \alpha_1 - \cos \alpha_2) \end{bmatrix}. \quad (\text{A9})$$

Here we used $m_1 = m_s + m_{1a}$, $m_2 = m_s + m_{2a}$, and $J = J_s + J_a$. Note that \mathbb{H} couples the translational and rotational motion of the articulated body. In the case of a single ellipsoid, \mathbb{H} is identically zero.

Further, $\mathbb{I}_{\text{couple}}$ is a 3×2 matrix that couples rigid-body motion with shape deformation:

$$\mathbb{I}_{\text{couple}} = \begin{bmatrix} -m_2 a \sin \alpha_1 & m_2 a \sin \alpha_2 \\ m_2 a \cos \alpha_1 & m_2 a \cos \alpha_2 \\ J + m_2 a^2(1 + \cos \alpha_1) & -J - m_2 a^2(1 + \cos \alpha_2) \end{bmatrix}. \quad (\text{A10})$$

Finally, $\mathbb{I}_{\text{shape}}$ is a 2×2 matrix associated with shape deformation:

$$\mathbb{I}_{\text{shape}} = \begin{bmatrix} J + m_2 a^2 & 0 \\ 0 & J + m_2 a^2 \end{bmatrix}. \quad (\text{A11})$$

The total linear and angular momenta P_1 , P_2 , and Π expressed in body frame are given by $P_1 = \partial T / \partial v_1$, $P_2 = \partial T / \partial v_2$, and $\Pi = \partial T / \partial \dot{\beta}$ to arrive at Eq. (3) of the main text:

$$\begin{bmatrix} P_1 \\ P_2 \\ \Pi \end{bmatrix} = \mathbb{I}_{\text{lock}} \begin{bmatrix} v_1 \\ v_2 \\ \dot{\beta} \end{bmatrix} + \mathbb{I}_{\text{couple}} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \quad (\text{A12})$$

APPENDIX B: PROXIMAL POLICY OPTIMIZATION (PPO) ALGORITHMS

We implement the clipped advantage proximal policy optimization (PPO) method proposed by [67] for our RL training. PPO maximizes a surrogate objective that clips off unwanted changes when the policy deviates too much from the policy of the previous cycle to ensure faster and more robust convergence. We refer readers to the original reference cited above as well as the OpenAI’s documentation of the PPO algorithm [70] and their baseline implementations [71] for a thorough explanation of the theory and details behind this method.

Our implementation can be separated into two parts. The main loop simulates the environment using action sequence a_t generated by the agent, and stores the observed rollouts for future updates (see Algorithm 1). Note that n_o and n_a are used to indicate the number of observable states and actions. Equations describing the fish-fluid interactions were integrated numerically using an adaptive time step, explicit RK45 method between each decision step of 0.1 unit of time. This choice of decision time step size limits the maximum rotation allowed for the fish head and tail to be 0.1 radian per step.

Parameters of the actor-critic networks of the RL agent are updated every N time steps for K epochs. Here the value of K is chosen to be 80 and the value of N is set to 4050, an integer multiple of the episode length 150. For simplicity, we assume our continuous action variables follow a multivariate normally distributed policy π_θ with mean value represented by a neural network parametrized by θ and constant diagonal covariance matrices, and the critic or value function $V_\phi(o_t)$ is also represented by a neural network with parameters ϕ . Specifically, both the mean policy and value function are implemented as feed-forward neural networks with two hidden layers and tanh activation functions. The sizes of the two hidden layers were fixed to 64 and 32, respectively. Each diagonal entry of the covariance matrix is set to 0.5². Finally, using the collected trajectories during the previous N time steps, the parameters θ, ϕ are updated according a total loss function $\mathcal{L}(\theta, \phi)$ via a back-propagating gradient-based optimizer (see Algorithm 2). Note that since we did not perform systematic hyperparameter tuning, readers might want to explore different values for better performance.

Algorithm 1. Environment simulation.

```

1: for time step  $t = 0, 1, \dots$  do
2:   if  $t = 0$  or episode terminates then
3:     store time step of episode termination,
4:     reset state  $s_t \sim P(s_0)$ 
5:     evaluate observation:  $o_t \sim o(s_t)$ 
6:   end if
7:   sample action from policy  $a_t \sim \pi_\theta(a_t|o_t)$ 
8:   evolve next state according to fish physics  $s_{t+1} \sim P(s_{t+1}|s_t, a_t)$ 
9:   evaluate next observation  $o_{t+1} \sim o(s_{t+1})$  and reward  $r_t \sim r(a_t, o_{t+1})$ 
10:  if  $t = 0$  or  $\text{mod}(t, N) \neq 0$  then
11:    append current action, observation, reward, and probability of sampling the action to assemble vectors
       $a_{N \times n_a}, o_{N \times n_o}, r_{N \times 1}$ , and  $\pi_{\theta_{\text{old}}}(a|o)_{N \times 1}$ 
12:  else
13:    update agent networks according to Algorithm II
14:  end if
15: end for

```

Algorithm 2. Updating the agent.

-
-
- 1: **for** update epoch number $\kappa = 0, 1, \dots, K$ **do**
 - 2: compute the truncated return using rewards $r_{N \times 1}$ and assemble into vector $R_{N \times 1}$
 - 3: estimate infinite-horizon return using $R_{N \times 1}$ and $V_T = V_\phi(o_T)$ if bootstrapping is desired [see Eq. (B2)]
 - 4: using $o_{N \times n_o}$ and value function V_ϕ , evaluate expected returns at each time step and store into $V_{N \times 1}$
 - 5: compute the advantage $A = R_{N \times 1} - V_{N \times 1}$ and normalize by its mean and variance if desired
 - 6: evaluate the probability of realizing $a_{N \times n_a}$ based on $o_{N \times n_o}$ for the policy π_θ , and store to $\pi_\theta(a|o)_{N \times 1}$
 - 7: compute the action-likelihood ratio: $\varrho_\theta = \frac{\pi_\theta(a|o)_{N \times 1}}{\pi_{\theta_{\text{old}}}(a|o)_{N \times 1}}$
 - 8: compute clipped surrogate loss function:
 $\mathcal{L}_{\text{clip}}(\theta) = \text{mean} [\min [\varrho_\theta \cdot A, \text{clip}(\varrho_\theta, 1 - \epsilon, 1 + \epsilon) \cdot A]]$
 - 9: compute the value-function loss: $\mathcal{L}_{\text{value}}(\phi) = 0.5 \times \text{mean} [(R_{N \times 1} - V_{N \times 1})^2]$
 - 10: compute the total loss: $\mathcal{L}(\theta, \phi) = -\mathcal{L}_{\text{clip}}(\theta) + \mathcal{L}_{\text{value}}(\phi) - \alpha \times \text{entropy} [\pi_\theta]$
 - 11: update parameters (θ, ϕ) to minimize the total loss using a gradient-based optimizer (e.g., Adam [72])
 - 12: **end for**
-
-

Another important side note is that since it is in general impossible to obtain an unrealized infinite horizon return $R_t = \sum_{t'}^\infty \gamma^{t'-t} r_{t'}$, we need to choose an appropriate estimator of this value based on finite length simulations. We can either simply truncate rewards after some step k by using

$$\hat{R}_t|_{\text{truncation}} = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots + \gamma^{k-1} r_{t+k}, \quad (\text{B1})$$

or we can use the trained value function (critic) to approximate the residual contribution to the return via k -step bootstrapping,

$$\hat{R}_t|_{\text{bootstrapping}} = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots + \gamma^k V_\phi(o_{t+k+1}). \quad (\text{B2})$$

We compared these two approaches for the direction-based task and observed that bootstrapping results in faster convergence and higher rewards in general [see Fig. 8(a)]. As a result, bootstrapping is used for all tasks depicted in the main text, where k was determined by the number of available




FIG. 8. Evolution of rewards during the training process. (a) Total rewards per episode achieved by policies trained to swim parallel to the x -axis in a driftless environment using bootstrapped (blue) and truncated (black) return estimates. Here solid lines indicate the median, and the shaded region shows the variation between 25th and 75th percentile for 24 runs of the learning algorithm. (b), (c) Total rewards obtained by policies trained to swim towards a given target, both of which adopt bootstrapped return estimates. Red in (b) represents naive policies trained in a driftless environment, while yellow in (c) represents policies trained in the presence of drift, with drift magnitude and direction supplied as additional observations to the policy. Again, lines and shaded regions indicate median and 25th to 75th percentile range, respectively.

future rewards. Namely, k decreased from 149 to zero as the number of time steps increased from 1 to 150 in each episode. In addition, we show the difference in training rewards and convergence speed between the naive policy and the drift-aware policy in Fig. 8(b). In general, the inclusion of more observations increased the time to convergence and variance in training rewards.

-
- [1] M. H. Dickinson, F. O. Lehmann, and S. P. Sane, Wing rotation and the aerodynamic basis of insect flight, *Science* **284**, 1954 (1999).
 - [2] E. Tytell, P. Holmes, and A. Cohen, Spikes alone do not behavior make: Why neuroscience needs biomechanics, *Curr. Opin. Neurobiol.* **21**, 816 (2011).
 - [3] J. Merel, M. Botvinick, and G. Wayne, Hierarchical motor control in mammals and machines, *Nat. Commun.* **10**, 1 (2019).
 - [4] M. S. Madhav and N. J. Cowan, The synergy between neuroscience and control theory: The nervous system as inspiration for hard control challenges, *Annu. Rev. Control Rob. Auton. Syst.* **3**, 243 (2020).
 - [5] K. H. Huang, M. B. Ahrens, T. W. Dunn, and F. Engert, Spinal projection neurons control turning behaviors in zebrafish, *Curr. Biol.* **23**, 1566 (2013).
 - [6] H. A. Burgess and M. Granato, Sensorimotor gating in larval zebrafish, *J. Neurosci.* **27**, 4984 (2007).
 - [7] M. Privat, S. A. Romano, T. Pietri, A. Jouary, J. Boulanger-Weill, N. Elbaz, A. Duchemin, D. Soares, and G. Sumbre, Sensorimotor transformations in the zebrafish auditory system, *Curr. Biol.* **29**, 4010 (2019).
 - [8] A. D. Marchese, C. D. Onal, and D. Rus, Autonomous soft robotic fish capable of escape maneuvers using fluidic elastomer actuators, *Soft Rob.* **1**, 75 (2014).
 - [9] G. V. Lauder, Fish locomotion: Recent advances and new directions, *Annu. Rev. Mar. Sci.* **7**, 521 (2015).
 - [10] M. Cianchetti, M. Follador, B. Mazzolai, P. Dario, and C. Laschi, Design and development of a soft robotic octopus arm exploiting embodied intelligence, in *2012 IEEE International Conference on Robotics and Automation* (IEEE, Piscataway, NJ, 2012), pp. 5271–5276.
 - [11] C. Laschi, M. Cianchetti, B. Mazzolai, L. Margheri, M. Follador, and P. Dario, Soft robot arm inspired by the octopus, *Adv. Rob.* **26**, 709 (2012).
 - [12] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, Cambridge, MA, 2006).
 - [13] R. Pfeifer, M. Lungarella, and F. Iida, Self-organization, embodiment, and biologically inspired robotics, *Science* **318**, 1088 (2007).
 - [14] W. M. Van Rees, M. Gazzola, and P. Koumoutsakos, Optimal shapes for anguilliform swimmers at intermediate Reynolds numbers, *J. Fluid Mech.* **722**, R3 (2013).
 - [15] J. C. Liao, D. N. Beal, G. V. Lauder, and M. S. Triantafyllou, Fish exploiting vortices decrease muscle activity, *Science* **302**, 1566 (2003).
 - [16] U. K. Müller, Fish'n flag, *Science* **302**, 1511 (2003).
 - [17] C. Eloy, On the best design for undulatory swimming, *J. Fluid Mech.* **717**, 48 (2013).
 - [18] S. Kern and P. Koumoutsakos, Simulations of optimized anguilliform swimming, *J. Exp. Biol.* **209**, 4841 (2006).
 - [19] R. Mittal, H. Dong, M. Bozkurtas, A. Loebbecke, and F. Najjar, Analysis of flying and swimming in nature using an immersed boundary method, in *Proceedings of the 36th AIAA Fluid Dynamics Conference and Exhibit* (AIAA, Washington, DC, 2006), Paper No. 2006–2867.
 - [20] J. D. Eldredge, Numerical simulations of undulatory swimming at moderate Reynolds number, *Bioinspiration Biomimetics* **1**, S19 (2006).
 - [21] M. Gazzola, M. Argentina, and L. Mahadevan, Gait and speed selection in slender inertial swimmers, *Proc. Natl. Acad. Sci. USA* **112**, 3874 (2015).
 - [22] E. D. Tytell, C. Y. Hsu, T. L. Williams, A. H. Cohen, and L. J. Fauci, Interactions between internal forces, body stiffness, and fluid environment in a neuromechanical model of lamprey swimming, *Proc. Natl. Acad. Sci. USA* **107**, 19832 (2010).

- [23] C. Hamlet, L. J. Fauci, and E. D. Tytell, The effect of intrinsic muscular nonlinearities on the energetics of locomotion in a computational model of an anguilliform swimmer, *J. Theor. Biol.* **385**, 119 (2015).
- [24] D. Guthrie, Role of vision in fish behaviour, in *The Behaviour of Teleost Fishes* (Springer, Berlin, 1986), pp. 75–113.
- [25] R. D. Fernald, Fish vision, in *Development of the Vertebrate Retina* (Springer, Berlin, 1989), pp. 247–265.
- [26] R. Douglas and M. Djamgoz, *The Visual System of Fish* (Springer, New York, 2012).
- [27] J. Engelmann, W. Hanke, J. Mogdans, and H. Bleckmann, Hydrodynamic stimuli and the fish lateral line, *Nature* **408**, 51 (2000).
- [28] L. Ristroph, J. C. Liao, and J. Zhang, Lateral Line Layout Correlates with the Differential Hydrodynamic Pressure on Swimming Fish, *Phys. Rev. Lett.* **114**, 018102 (2015).
- [29] B. Colvert and E. Kanso, Fishlike rheotaxis, *J. Fluid Mech.* **793**, 656 (2016).
- [30] B. Colvert, G. Liu, H. Dong, and E. Kanso, How can a source be located by responding to local information in its hydrodynamic trail?, in *IEEE 56th Annual Conference on Decision and Control (CDC)* (IEEE, Piscataway, NJ, 2017), pp. 2756–2761.
- [31] J. C. Montgomery, C. F. Baker, and A. G. Carton, The lateral line can mediate rheotaxis in fish, *Nature* **389**, 960 (1997).
- [32] A. Filella, F. Nadal, C. Sire, E. Kanso, and C. Eloy, Model of Collective Fish Behavior with Hydrodynamic Interactions, *Phys. Rev. Lett.* **120**, 198101 (2018).
- [33] L. Li, M. Nagy, J. M. Graving, J. Bak-Coleman, G. Xie, and I. D. Couzin, Vortex phase matching as a strategy for schooling in robots and in fish, *Nat. Commun.* **11**, 1 (2020).
- [34] B. L. Partridge and T. J. Pitcher, The sensory basis of fish schools: Relative roles of lateral line and vision, *J. Comp. Physiol.* **135**, 315 (1980).
- [35] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin, Inferring the structure and dynamics of interactions in schooling fish, *Proc. Natl. Acad. Sci. USA* **108**, 18720 (2011).
- [36] M. Gazzola, A. A. Tchieu, D. Alexeev, A. de Brauer, and P. Koumoutsakos, Learning to school in the presence of hydrodynamic interactions, *J. Fluid Mech.* **789**, 726 (2016).
- [37] K. Gustavsson, L. Biferale, A. Celani, and S. Colabrese, Finding efficient swimming strategies in a three-dimensional chaotic flow by reinforcement learning, *Eur. Phys. J. E* **40**, 110 (2017).
- [38] S. Colabrese, K. Gustavsson, A. Celani, and L. Biferale, Flow Navigation by Smart Microswimmers via Reinforcement Learning, *Phys. Rev. Lett.* **118**, 158004 (2017).
- [39] S. Verma, G. Novati, and P. Koumoutsakos, Efficient collective swimming by harnessing vortices through deep reinforcement learning, *Proc. Natl. Acad. Sci. USA* **115**, 5849 (2018).
- [40] B. Colvert, M. Alsalman, and E. Kanso, Classifying vortex wakes using neural networks, *Bioinspiration Biomimetics* **13**, 025003 (2018).
- [41] M. Alsalman, B. Colvert, and E. Kanso, Training bioinspired sensors to classify flows, *Bioinspiration Biomimetics* **14**, 016009 (2018).
- [42] P. Weber, G. Arampatzis, G. Novati, S. Verma, C. Papadimitriou, and P. Koumoutsakos, Optimal flow sensing for schooling swimmers, *Biomimetics* **5**, 10 (2020).
- [43] M. Gazzola, B. Hejazialhosseini, and P. Koumoutsakos, Reinforcement learning and wavelet adapted vortex methods for simulations of self-propelled swimmers, *SIAM J. Sci. Comput.* **36**, B622 (2014).
- [44] E. Kanso and J. E. Marsden, Optimal motion of an articulated body in a perfect fluid, in *Proceedings of the 44th IEEE Conference on Decision and Control* (IEEE, Piscataway, NJ, 2005), pp. 2511–2516.
- [45] E. Kanso, J. E. Marsden, C. W. Rowley, and J. B. Melli-Huber, Locomotion of articulated bodies in a perfect fluid, *J. Nonlinear Sci.* **15**, 255 (2005).
- [46] M. H. Dickinson, C. T. Farley, R. J. Full, M. Koehl, R. Kram, and S. Lehman, How animals move: An integrative view, *Science* **288**, 100 (2000).
- [47] T. Degris, P. M. Pilarski, and R. S. Sutton, Model-free reinforcement learning with continuous action in practice, in *Proceedings of the 2012 American Control Conference (ACC)* (IEEE, Piscataway, NJ, 2012), pp. 2177–2182.
- [48] A. M. Haith and J. W. Krakauer, Model-based and model-free mechanisms of human motor learning, in *Progress in Motor Control*, edited by M. J. Richardson, M. A. Riley, and K. Shockley (Springer, New York, 2013), pp. 1–21.

- [49] A. C. H. Tsang, P. W. Tong, S. Nallan, and O. S. Pak, Self-learning how to swim at low Reynolds number, *Phys. Rev. Fluids* **5**, 074101 (2020).
- [50] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [51] J. D. Eldredge, Numerical simulation of the fluid dynamics of 2D rigid body motion with the vortex particle method, *J. Comput. Phys.* **221**, 626 (2007).
- [52] J. B. Melli, C. W. Rowley, and D. S. Rufat, Motion planning for an articulated body in a perfect planar fluid, *SIAM J. Appl. Dyn. Syst.* **5**, 650 (2006).
- [53] S. Nair and E. Kanso, Hydrodynamically coupled rigid bodies, *J. Fluid Mech.* **592**, 393 (2007).
- [54] R. L. Hatton and H. Choset, Connection vector fields and optimized coordinates for swimming systems at low and high Reynolds numbers, in *Proceedings of the ASME 2010 Dynamic Systems and Control Conference* (American Society of Mechanical Engineers, New York, 2010), pp. 817–824.
- [55] K. A. Morgansen, V. Duidam, R. J. Mason, J. W. Burdick, and R. M. Murray, Nonlinear control methods for planar carangiform robot fish locomotion, in *Proceedings of the 2001 IEEE International Conference on Robotics and Automation* (IEEE, Piscataway, NJ, 2001), Vol. 1, pp. 427–434.
- [56] A. M. Bloch, P. Krishnaprasad, J. E. Marsden, and R. M. Murray, Nonholonomic mechanical systems with symmetry, *Arch. Ration. Mech. Anal.* **136**, 21 (1996).
- [57] R. M. Murray, Z. Li, S. S. Sastry, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation* (CRC Press, Boca Raton, FL, 1994).
- [58] S. D. Kelly, P. Pujari, and H. Xiong, Geometric mechanics, dynamics, and control of fishlike swimming in a planar ideal fluid, in *Natural Locomotion in Fluids and on Surfaces* (Springer, Berlin, 2012), pp. 101–116.
- [59] H. Lamb, *Hydrodynamics* (Dover, New York, 1945).
- [60] R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Mach. Learn.* **8**, 229 (1992).
- [61] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, Policy gradient methods for reinforcement learning with function approximation, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, USA, 2000), pp. 1057–1063.
- [62] S. M. Kakade, A natural policy gradient, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, USA, 2002), pp. 1531–1538.
- [63] J. Peters and S. Schaal, Natural actor-critic, *Neurocomputing* **71**, 1180 (2008).
- [64] J. Peters and S. Schaal, Policy gradient methods for robotics, in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, Piscataway, NJ, 2006), pp. 2219–2225.
- [65] V. R. Konda and J. N. Tsitsiklis, Actor-critic algorithms, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, USA, 2000), pp. 1008–1014.
- [66] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, in *Proceedings of the 35th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by J. Dy and A. Krause (PMLR, Stockholmsmässan, Stockholm Sweden, 2018).
- [67] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- [68] <https://github.com/mjysh/RL3linkFish>
- [69] N. E. Leonard, Stability of a bottom-heavy underwater vehicle, *Automatica* **33**, 331 (1997).
- [70] <https://spinningup.openai.com/en/latest/algorithms/ppo.html>
- [71] <https://github.com/openai/baselines/tree/master/baselines/ppo2>
- [72] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *3rd International Conference on Learning Representations*, edited by Y. Bengio and Y. LeCun (ICLR, San Diego, CA, USA, 2015).

Correction: The previously published order of authors was presented incorrectly and has been fixed, and the affiliations are now identified by superscript numbers.