

Control Problems Solved by a Fish's Body and Brain: A Review

J. Edward Colgate and Kevin M. Lynch

Abstract— The bodies and brains of fish have evolved to achieve control objectives beyond the capabilities of current underwater vehicles. One route toward designing underwater vehicles with similar capabilities is to better understand fish physiological design and control strategies. This paper has two objectives: (1) to review clues to artificial swimmer design taken from fish physiology, and (2) to formalize and review the control problems which must be solved by a robot fish. The goal is to exploit fish locomotion principles to address the truly difficult control challenges of station keeping under large perturbations, rapid maneuvering, power-efficient endurance swimming, and trajectory planning and tracking. The design and control of biomimetic swimming machines meeting these challenges will require state-of-the-art engineering and biology.

Keywords— Robot fish design, swimming modes, controllability, station keeping, trajectory tracking, motion primitives

I. INTRODUCTION

The bodies and brains of fish have evolved to achieve control objectives beyond the capabilities of current underwater vehicles. Some fish are highly maneuverable, some are power-efficient endurance swimmers, some are accomplished at station keeping in the presence of significant perturbations, and some combine these capabilities. One route toward designing underwater vehicles with similar capabilities is to better understand fish physiological design and control strategies.

This paper has two objectives: (1) to review clues to artificial swimmer design taken from fish physiology, and (2) to formalize and review the control problems which must be solved by a robot fish. Design encompasses issues such as morphology, passive mechanical properties, actuator selection and placement, and sensor selection and placement. Control encompasses issues such as planning and executing goal-directed fin and body movements and incorporating sensory feedback. Because there have been many more studies on the control advantages offered by fish mechanical design and behavior than on neural feedback control of fish locomotion, our review of the control problems to be solved by robot fish focuses on control-theoretic approaches to robot fish locomotion.

As a first step, it is necessary to understand context and goals. A robot built to locomote in littoral zones is likely quite different from one built to swim in deep water. Similarly, a robot built to swim efficiently for long distances will be different from one built to perform rapid maneuvers.

To limit the scope of this review, we first take note of the fact that it is relatively straightforward, with existing technology, to build robotic fish that look life-like as they swim through the water. Among the most convincing of these are the artificial sea bream and coelacanth (an extinct fish) developed by Mitsubishi Heavy Industries. Our goal is not to build life-like robotic fish, but to exploit fish locomotion principles to address the truly dif-

ficult control challenges of station keeping under large perturbations, rapid maneuvering, power-efficient endurance swimming, and trajectory planning and tracking. The design and control of biomimetic swimming machines meeting these challenges will require state-of-the-art engineering and biology.

Section II reviews fish swimming modes, the basic hydrodynamics of swimming, and the effect of design choices such as fin placement and stiffness on maneuverability and stability. Section III formulates the control problems which must be solved by a robot fish, taking consideration of the possible underactuated nature of the fish, and reviews proposed solutions to these problems.

II. CONSIDERATIONS FOR THE DESIGN OF BIOMIMETIC SWIMMING MACHINES

A. Morphology and Swimming Modes

A basic consideration for the design of swimming machines is the design of propulsors: their shape, their location on the robot, their mechanical properties (e.g., inertia and stiffness), and their pattern of movement. Another important consideration is the overall shape of the robot. Fish provide useful illustrations of propulsor design, swimming modes, and body shape (morphology). Not surprisingly, they also reveal that these are tightly interrelated [89]. Sfakiotakis *et al.* [105] provide an excellent review of these factors, including the classification scheme illustrated in Figure 1. Their classification is based on two main factors: (1) the extent to which propulsion is based on undulatory motion versus oscillatory motion, and (2) the body structures or fins that are most active in generating thrust. In order to illustrate these factors and review a few of the most relevant swimming modes and fish morphologies, we will discuss several examples below. For a more complete discussion, the reader is referred to [105].

Fish that generate thrust principally via body and/or caudal tailfin motions are known collectively as BCF swimmers (Figure 1a). For example, in *anguilliform* locomotion, the entire body is the propulsor. Anguilliform swimming, which is typical of eels, lampreys and the invertebrate amphioxus (lancelets), requires a long, slender, and highly flexible body. Thrust is produced by undulation, i.e., passing a transverse wave from head to tail. Thus, one useful feature of anguilliform swimming is the ability to swim backwards, which simply requires that the wave be passed from tail to head [20]. Robotic anguilliforms have been described in [7], [83], and have proved useful in the study of artificial neural network control as well as path planning for underactuated systems (see Section III). *Carangiform* locomotion also involves undulation of the entire body, but with amplitude growing toward the tail. Typical carangiform swimmers (jacks, mackerel, snapper) have a narrow peduncle and a




Fig. 1. Swimming modes associated with (a) BCF propulsion and (b) MPF propulsion. Shaded areas contribute to thrust generation. Modified from [105].

tall, forked caudal fin. These are among the swiftest of swimmers. The fastest of the carangiforms are often placed in another category, known as *thunniform* swimmers. These fish, including tuna and some sharks, have very low-drag body shapes, narrow peduncles, and tall, lunate (crescent-shaped) caudal fins. Not surprisingly, most studies of efficient fin-based swimming have focused on the carangiforms and thunniforms. Mason and Burdick [79] and Saimek and Li [102] describe robotic testbeds for the study of carangiform swimming, while the MIT RoboTuna is an example of a robotic thunniform [110]. *Sub-carangiform* locomotion is intermediate between carangiform and anguilliform.

Although the caudal fin is arguably the most successful of aquatic propulsors (certainly based on number of species), there are many types of fish (Figure 1b), known as MPF swimmers, that generate thrust using principally median (e.g., dorsal and anal) and paired (e.g., pectoral) fins. Among these, a particularly important class is the *labriform* swimmers that make use of pectoral fin oscillation and almost no body motion. These include the wrasses, angelfish and surperches. As will be discussed later, the labriform swimmers may generate either drag-based or lift-based locomotion. While there is evidence that labriform swimmers are less efficient than carangiform swimmers at cruising speeds [105], labriform swimming is thought to be more efficient at slow speeds and considerably more maneuverable as well [17], [18]. Maneuverability stems from the fact that the two pectoral fins may be controlled independently of one another, and also from the fact that pectoral fins may produce reverse thrust. Labriforms also appear to be proficient at station-keeping in currents, which may relate to the absence of caudal fin oscillations, reducing lateral recoil and sideslip [45]. Kato [55], [56] has performed extensive testing and control development for a labriform robot known as “Bass II.”

The *ostraciiform* swimmers, such as cowfish and boxfish, make use of both median and paired fins to swim. These fish are also characterized by a rigid carapace that encases about three-fourths of the body length. Although classically known as slow swimmers, a recent study points to high maximum speeds, high endurance, and a high level of dynamic stability, all features that would be desirable in an AUV [45].

A final swimming mode that will be mentioned here is that of the *gymnotiforms*, especially the American knifefishes. These fish possess a ventral ribbon fin that runs most of the length of the body. They have pectoral fins and either a vestigial caudal fin or none at all. Thrust is produced by passing waves down the ribbon fin while the body remains relatively rigid. These fish also possess an elaborate electrosensory system with receptors spread across their bodies. The gymnotiforms are remarkably maneuverable swimmers, able to reverse swimming direction suddenly and also roll their bodies in order to capture prey [77]. This maneuverability stems from the fact that the ribbon fin is a propulsor having a high number of actively controlled inputs. By properly coordinating these inputs, exquisite control of the thrust vector is possible [105], [19].

B. Non-dimensional parameters and key metrics

To facilitate the discussion of swimming, especially comparisons across different fish, several non-dimensional parameters and metrics are commonly used. **The propulsive efficiency, or Froude number**, is defined as

$$\eta_P = \frac{P_E}{P_P}, \quad (1)$$

where $P_E = UT_A$ is the useful propulsive power, the product of the time-averaged thrust T_A and swimming speed U , and P_P is the time-averaged power expended by the fish [14]. Propulsive

efficiencies reported in the literature vary over a considerable range, from 90% for a carangiform swimmer [90] to only 16% for drag-based labriform locomotion [18].

The reduced frequency or Strouhal number relates oscillation frequency f to swimming speed U , and is considered a measure of the importance of unsteady hydrodynamic effects [105], [34]. It is defined as

$$\sigma = 2\pi \frac{fL}{U} \quad (2)$$

where L is a characteristic length, usually the chord length of the moving fin. If σ is greater than 0.1, unsteady effects are considered important. This is the case for most fish. Triantafyllou [108] introduces a closely related measure where the length scale A derives from the peak-to-peak lateral excursion of the tail fin at its junction with the body. Following [110], we will simply call this the Strouhal number:

$$St = \frac{fA}{U}. \quad (3)$$

Triantafyllou and Triantafyllou [110] argue that a Strouhal number in the range 0.25-0.35 is indicative of efficient swimming.

The Reynolds number is approximated as

$$Re = \frac{UL}{\nu}, \quad (4)$$

where ν is the kinematic viscosity of water. For most fish, Re is in the range 10^3 to 5×10^6 , for which viscous forces are negligible, but added mass, pressure drag and lift are all important [105].

C. Mechanics and models of swimming

Fish generate thrust by transferring momentum to the surrounding fluid. While the detailed mechanisms of this momentum transfer are typically quite complex, it is often possible to develop greatly simplified models that are competent to support controller design. Not all energy generated by a fish produces useful thrust, however; some of it is dissipated to the surrounding fluid via mechanisms which may also be quite complex. But here again, simple models often prove quite useful. In the sequel, a brief review is given of both thrust and dissipation mechanisms, followed by commentary on model development for real-time control. Much more detailed discussions of biohydrodynamics are given elsewhere in this volume [109], [84].

When a fish moves its body or fins relative to the surrounding fluid, fluid is displaced and reaction forces arise. The net force and moment that accelerate (or decelerate) the fish's body are given simply by surface integrals of the vector force per unit area taken over the body. In principle, this force distribution as well as the surrounding flow field may be computed by solution of the Navier-Stokes equation. In practice, this requires computational fluid dynamics codes [84], a subject beyond the scope of this review. Many useful insights, however, can be obtained by making simplifications and describing the fluid mechanics in terms of constituent effects.

Classically, an important simplification is the assumption of steady or quasi-steady flow. The forces on a foil immersed in a steady flow may be described as lift and drag. Lift is defined as the force acting perpendicular to the direction of motion, while

drag acts parallel to the direction of motion. At high Reynolds number, the lift per unit area, L , and drag per unit area, D , take the form

$$L = \frac{\rho C_L U^2}{2}, \quad D = \frac{\rho C_D U^2}{2}. \quad (5)$$

Here, C_L and C_D are nondimensional lift and drag coefficients, respectively. They depend on the shape of the foil, the Reynolds number, and the angle of attack. If the dependence on angle of attack is known for a particular foil, then these equations may be used in the *blade-element* method to compute thrust and drag. This is accomplished by specifying the foil's motion while discretizing time and treating the foil's surface as a collection of blade elements. At each point in time, lift and drag are computed for each blade element, then summed to compute total thrust, moment and drag. A good example of this technique applied to pectoral fins may be found in [17]. Here the results of the blade element analysis are averaged over a fin beat cycle to study Froude efficiency.

There are a number of limitations to the blade element method, the most obvious of which is that acceleration reaction forces — those stemming from fluid acceleration — are neglected. Acceleration reaction may be captured fairly simply by endowing the foil with “added mass” to account for the fluid inertia that moves along with it [17], [115]. Unsteady effects, however, are more difficult to capture and ultimately limit the power of this technique.

Dickinson has given an excellent review of unsteady mechanisms [34]. His analysis employs a more sophisticated view of hydrodynamics, focusing especially on circulation and vorticity. The reader is referred to his paper for a complete discussion; however, a few key ideas will be reviewed here. First, the lift on a foil is proportional to the circulation around it. Yet, Kelvin's Law tells us that the net circulation within a fluid system is constant. In the case of a foil starting from rest, the net circulation must be zero. This begs the question: how can lift be generated? The answer is that, as the foil begins to move and circulation develops about it, a *starting vortex* is shed from its trailing edge. As the foil departs from the starting vortex, the effect of that vortex diminishes, while the bound circulation remains, producing lift. When stopping, the bound circulation is shed as a stopping vortex, equal and opposite in strength to the original starting vortex, and lift once again drops to zero. This picture of lift generation is particularly valuable when considering oscillating fins, which may be viewed as continually starting and stopping. If the reduced frequency is sufficiently high, we may surmise that the starting and stopping vortices never move very far from one another and must be considered in any computation of thrust. This phenomenon was first described by Wagner, who found that the starting vortex tends to counteract circulation, reducing lift, so long as it remains within a few chord lengths of the foil.

Of even greater importance for fish, however, is the unsteady phenomenon of *delayed stall*. Delayed stall occurs during the impulsive start of a foil at a high angle of attack. As the foil begins to move, a starting vortex is shed from the trailing edge, but instead of bound circulation growing around the foil, an attached vortex develops at the leading edge. This vortex creates a region of low pressure on the top surface of the foil resulting in

lift that may be much greater than that arising from steady state circulation. Dickinson suggests that this effect may be of enormous importance in swimming, and may explain why many fish use high angles of attack. Unfortunately, delayed stall is very difficult to capture in a simple model.

As a fish flaps its fins, the leading edge vortex created by delayed stall will eventually be shed, and a new, counter-rotating vortex will develop and be shed over the next half-cycle of oscillation. This process leads to a wake of alternating vortices known as a Karman street. This wake structure has been described in detail in [34], [14], [64].

A final point to make with respect to the discussion of unsteady effects, especially at high angles of attack, is that lift and drag become closely related. For instance, in delayed stall the force that is developed is roughly normal to the foil surface. But, owing to the high angle of attack, the projection of this force onto the direction of motion, i.e. the drag, is quite significant. This is an example of *pressure drag*, which can be much more important than the direct effect of viscous shear, known as *skin friction*. Skin friction is inversely proportional to the square root of Reynolds number, and tends to be unimportant for most fish. A third component of lift, *induced drag*, results from the rearward component of lift arising from bound circulation, and occurs even at small angles of attack. The drag coefficient used in steady analysis would normally correspond to pressure drag and induced drag.

A challenge for the designer of biomimetic AUVs is the development of hydrodynamic models that may be used for motion planning and real-time control. Here, simplicity is clearly a virtue. It is not surprising, therefore, that most models developed for such purposes make various simplifying assumptions. For example, in their work on robotic anguilliforms, McIsaac and Ostrowski [83] use what is essentially a blade element approach in which the only forces arise from pressure drag.

Perhaps the best-known model of swimming is Lighthill's *elongated body theory* (EBT), which has been used to study anguilliform and carangiform propulsion. EBT makes use of linear, inviscid aerodynamic theory (see, for instance, [42]). Although strictly speaking an unsteady theory because it assumes sinusoidal pitching and heaving motions, EBT has been used principally to study the mechanics of steady swimming. Of greater interest for AUV control is the related model developed by Harper, Berkemeier and Grace [49] which comprises a set of low-order, linear ordinary differential equations (ODEs). Such equations are well-suited for the development of feedback controllers. Harper's model assumes that the flow is inviscid (a reasonable assumption at high Reynolds numbers) and incompressible, as well as irrotational outside a region defined by the foil and its wake. Although the flow is taken to be inviscid, it must obey the Kutta-Joukowski condition that velocity be finite at the sharp trailing edge of the foil, a condition that arises from viscosity. Moreover, while the flow is irrotational outside the foil and wake, the model does allow for circulation about the foil and vorticity (essentially a starting vortex) in the wake. The point is that these assumptions are not terribly restrictive. Unfortunately, in order to linearize the equations, it is also necessary to assume small foil displacements, including a small angle of attack. While doing so may be considered restrictive, it al-

lows the wake effect to be described in the frequency domain via the Theodorsen function [42]. Harper's model approximates the Theodorsen function with a low-order linear filter, leading directly to a set of ODEs for the lift and moment acting on the foil. This model was used to optimize the design and stability characteristics of an oscillating foil propulsor.

Other sets of simplifying assumptions may also prove quite useful in developing models. For instance, Saimek and Li [102] assume that the flow is inviscid, incompressible, and completely irrotational (including circulation about the foil and in the wake). They also assume that the flow is at rest infinitely far from the foil, and that the foil is neutrally buoyant. This may seem like an overly restrictive set of assumptions since lift is ignored, but added mass effects remain, and this model may form the basis of an effective control policy. **Under this set of assumptions the Lagrangian of the dynamic system (foil and fluid) is simply the total kinetic energy, which can be written**

$$L(q, \dot{q}) = \frac{1}{2} \dot{q}^T M(q) \dot{q},$$

where $q \in \mathbb{R}^n$ is the vector of generalized coordinates describing the foil's configuration, and $M(q)$ is the $n \times n$ symmetric positive-definite inertia matrix due to the foil's mass and the added mass of the fluid. For conservative systems of this form, the equations of motion can be written

$$M(q)\ddot{q} + C(q, \dot{q})\dot{q} = T(q)u, \quad (6)$$

where the Coriolis matrix $C(q, \dot{q})$ is linear in \dot{q} , $u \in \mathbb{R}^m$ is the control vector, and $T(q)$ is an $n \times m$ matrix indicating how the controls act on the generalized coordinates. There is an extensive body of robotics literature that applies to the control of systems in this form, as discussed in Section III.

Kelly *et al.* [58], [60] develop a model having some of the features of both those described above. Their model, which applies to carangiform swimmers, addresses the mechanics of not only the foil (caudal fin), but the body as well. They develop a set of reduced Euler-Lagrange equations to describe the interaction of the body with the surrounding fluid, which is treated as incompressible and inviscid, but containing a point vortex termed a *substitution vortex*. The substitution vortex represents the circulation developed around the caudal fin. The validity of this approach stems from the morphology of carangiforms: because of the long, narrow peduncle, the caudal tailfin is located several chord lengths away from the main body. At this distance, the flow field is largely independent of fin shape, and may be treated as though it originated from a point vortex. The strength of this vortex is computed, however, using an unsteady model similar to Harper's but not including the wake effect.

The models described above have great appeal for roboticists since they are based on the hydrodynamics of swimming, yet are low order and tractable for control analysis. It should be pointed out, however, that there exist other approaches to the development of models that circumvent the difficulties of theoretical hydrodynamics and, possibly, the need to make restrictive assumptions. This topic is further discussed in Section III-C.

D. Stability and Maneuverability

Open- and closed-loop stability are important considerations for the design of any control system, including one for aquatic

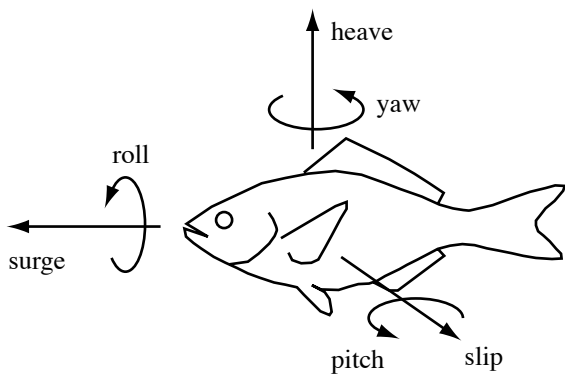


Fig. 2. Motions in a body frame.

locomotion. Closed-loop stability is clearly essential, but the same cannot be said for open-loop stability. It is well established that there exists a tradeoff between stability and maneuverability [39], [116], [118], and that many fish are open loop unstable. For instance, most dead fish tend to float belly-up, suggesting that the center of buoyancy lies below the center of mass. Indeed this is true [89], although there are counter-examples [117]. Thus, most fish are inverted pendulums, and must use feedback control to remain upright. Also, Weihs points out that BCF swimmers, because they are essentially pushed from behind, tend to be unstable in yaw [118] (Figure 2). Pitch instabilities are also common because the center of buoyancy is typically located fore or aft of the center of mass [117]. Hydrostatic pitching moments must be counteracted with propulsors, which may explain why most fish are observed to beat their paired fins even at rest.

The most common benefit attributed to open loop instability is high maneuverability. Indeed, many fish are quite maneuverable. Bandyopadhyay [12] compares normal acceleration to turning radius (normalized by body length) and finds that fish are able to execute comparable turns at much lower normal acceleration than two small underwater vehicles. It is not clear, however, whether this should be attributed to instability, or simply the fish's ability to generate turning moments at low speed, which foils such as rudders cannot. Anderson and Chhabra [4] demonstrate dramatic increases in yaw rate during turning maneuvers for a robotic tuna as compared to a conventional underwater vehicle.

While open loop instability may be a good thing, the problem remains of how to achieve closed loop stability, which is essential for activities such as cruising and station holding. Webb describes two basic strategies that fish use: trimming forces and powered correction forces [116]. The former arise from steering the flow over control surfaces in order to develop corrective forces. In a recent study of the ostraciiform trunkfish, however, Bartol *et al.* concluded that the unique shape of the bony carapace led to trimming-based stability in both pitch and yaw [15]. Nonetheless, for most fish trimming forces arise from the fins, leading to a design tradeoff: larger fins produce stronger trimming forces and enhance stability, but also increase drag and reduce maneuverability [4]. Powered correction forces arise from the flapping of fins. While powered correction is energetically

more costly than trimming, the latter becomes ineffective at low speeds requiring that powered correction be used for stable hovering [16]. To the best of the authors' knowledge, no robotic studies have yet addressed the use of paired fins for stabilization at either low or high speeds.

Indeed, very little work has been done that would indicate what the open-loop stability properties of a robotic swimmer ought to be, or how this would relate to morphology and swimming mode. Anderson and Chhabra [4] make use of hydrodynamic coefficients to study the stability of the Vorticity Controlled Unmanned Undersea Vehicle (VCUUV), a direct descendant of the RoboTuna [110]. Hydrodynamic coefficients are generalizations of the standard lift and drag coefficients. They relate the six-vector of forces and torques acting on the rigid body coordinates to the six-vector of velocities and angular velocities, as well as perturbations of body coordinates such as caudal fin angle. The use of these coefficients is a powerful technique permitting not only analysis, but design, including the sizing and placement of fins.

As a final point, it is important to remember that maneuverability is very much related to swimming mode and body morphology. For instance, as described earlier, the gymnotids are highly maneuverable in roll, an ability stemming from their compressed body shape and long ribbon fin. Similarly, in an effort to increase yawing maneuverability, Bandyopadhyay [11] added a long dorsal fin capable of controlled camber to a small underwater vehicle. Labriforms are particularly interesting from a maneuverability standpoint. Bandyopadhyay [12] points out that a number of fish which are both fast and agile make use of pectoral fins, and Kato [55], [56] cites maneuverability at low speeds as a rationale for the study of mechanical pectoral fins. Westneat *et al.* [119] (this volume) give a detailed review of pectoral fin morphology and kinematics as they relate to maneuverability. Lauder and Drucker [63] (this volume) summarize a set of lessons for the placement of fins on a biomimetic AUV.

E. Actuation and Passive Mechanical Properties

Fish use muscle to actuate their propulsors. As muscle is not currently an alternative for artificial systems, a variety of other actuators have been explored. Most commonly, electric motors are employed [14], [79], [102], [55], but a difficulty with motors is that they must generally be used along with fairly elaborate transmission systems to amplify torque and convey it to the propulsors [14]. Alternatives that arguably provide configuration advantages, if not efficiency advantages, include shape memory alloy actuators [96], piezoelectric unimorph actuators [22], and ICPF actuators [47].

A more direct cue can be taken from the manner in which fish make use of passive mechanical properties — most notably springiness — to enhance the effectiveness of their muscles. For instance, it stands to reason that, in generating oscillatory motions, the need for muscles to perform negative work will be minimized if the system is driven at its natural frequency. Springs, including the muscles and tendons themselves, collagen fibers, skin and other body structures [99], [53], play an important role in defining the natural frequency. The significance of this concept to robot designers was demonstrated by Harper, Berkemeier and Grace [49] who computed a 33% reduction in

energy costs for heaving motions when driving a hydrofoil via a properly tuned series stiffness. An important issue, however, is that different swimming speeds require different stiffnesses in order to match body natural frequency to tailbeat frequency. McHenry, Pell and Long [82] built vinyl models of the pumpkinseed sunfish to investigate the effect of body stiffness on undulatory waveforms. One interesting conclusion is that the live pumpkinseed must actively increase body stiffness to double the level of passive stiffness in order to achieve their swimming speeds. In a related study, Long and Nipper [74] explored the use of muscles to modulate body stiffness. They showed that the muscles of a largemouth bass are capable of modulating stiffness. MacIver [76] points out that gymnotids have a special body-stiffening adaptation, intra-muscular bones (although this adaptation serves to maintain a rigid body shape during prey capture rather than modulate natural frequency). Using muscles to stiffen the body requires that the muscles produce negative work, which seems to invalidate the primary rationale for tuning body stiffness — energetic efficiency. There may, however, be other energetic advantages to a properly tuned stiffness that outweigh the cost of negative work. As an example, Barrett *et al.* [14] use the RoboTuna to demonstrate significant drag reduction associated with undulation. In a robot, stiffness modulation does not need to be accomplished actively, potentially allowing both benefits: reduced negative work and reduced drag. In studies of a robotic dolphin with an actuated joint at the peduncle and a passive compliant joint at the caudal fin, Nakashima and Ono [94] found correlations of speed, tailbeat frequency and reduced frequency with stiffness of the compliant joint, providing a possible basis for performance optimization.

In his extensive studies of animal locomotion, Alexander has pointed out that springs have several important functions beyond energy efficiency [3]. These include, for instance, energy storage for catapulting, shock absorption, and acting as muscle antagonists. In her review of springs in swimming, Pabst [99] has highlighted the latter function. As an example, the skin of a lemon shark is reinforced with helically wound collagen fibers, producing a distinctly nonlinear (hardening) spring that acts in parallel with the musculature. As the shark’s body is flexed side to side, the outside muscles may get stretched to a point where they are ineffective at generating force. But at this point the springs are highly stretched, and may serve to accelerate unbending.

Numerous other mechanical aspects of fish, such as body shape, mucous layers, and scales, are also known to affect efficiency and speed, but are beyond the scope of this review.

F. Sensing

A final set of design considerations, prior to the development of control mechanisms, is the nature and allocation of sensors. Fish have the usual array of olfactory, acoustical, visual and tactile sensors, adapted of course to their aqueous environment. Fish also have an inner ear structured much like our own which is critical for maintaining orientation. Many fish also have specialized sensory apparatus, such as the electrosense of the gymnotids, which is discussed in detail by MacIver elsewhere in this volume [76].

But perhaps most interesting to designers of robotic swim-

mers is the *lateral line* system common to all fish. This system consists of a set of hair cell mechanoreceptors known as neuromasts. The superficial neuromasts extend outward from the epidermis into the flow, while canal neuromasts lie beneath the epidermis and communicate with the flow via a set of pores. It has been shown that canal neuromasts respond better to the acceleration of the water, while superficial neuromasts respond to the velocity of the flow across the fish’s body [85].

There is considerable evidence that the lateral line system provides fish with a “distance touch” capability that is used to detect other animals (e.g. for schooling or prey capture) and obstacles [31], [95]. Montgomery *et al.* [85] have shown that the superficial neuromasts play an important role in station keeping (rheotaxis). Several varieties of fish were able to orient to current at much lower flow rates with the superficial lateral line system intact than with it blocked. The canal neuromasts seemed to play no role in rheotaxis, as one might expect, given their greater sensitivity to water accelerations.

Less is known about the intriguing possibility that the lateral line may be used in the closed-loop control of swimming motions, as has been suggested by Triantafyllou [110]. Li and Saimek [72] have developed a Kalman filter based estimation scheme which recovers the hydrodynamic potential (assuming inviscid, incompressible flow and a wake comprised of point vortices) from a set of pressure measurements along a fish’s body. Such a method is likely to be an important precursor to the development of effective closed-loop controllers. On the hardware side, Fan *et al.* [38] have developed flow sensors reminiscent of superficial neuromasts using silicon micromachining. A cautionary note, however, is that there is little biological evidence for the role of the lateral line in swimming performance. Indeed, Dijkgraaf [35] reports that the swimming performance of blinded minnows was unchanged after complete elimination of the lateral line system. If the labyrinth was removed instead, leaving the lateral line intact, the minnows were no longer able to swim in a normal manner.

III. CONSIDERATIONS FOR THE CONTROL OF BIOMIMETIC SWIMMING MACHINES

Control problems for a robot fish include the following:

1. Station keeping in the presence of disturbances, such as in a littoral zone.
 2. Trajectory planning for point-to-point motions, possibly in the presence of obstacles.
 3. Tracking a planned trajectory.
 4. Power efficient swimming, for long battery life and “quiet” wake signatures.
 5. Fast starts and high maneuverability, e.g., the ability to turn 180 degrees within one body length, even at high speeds.
- These issues are naturally of interest for any underwater vehicle, biomimetic or otherwise. In this section we take inspiration from fish while phrasing the problems in modern control terms.

A. Modeling and Controllability

Usually the first step to designing controllers to address these issues is to write a dynamic model of the system. Generically, such a model takes the form

$$\dot{x} = f(x, u), \quad (7)$$




Fig. 3. A planar model of a robot fish.

where x is the state of the system, u is the control vector, and time-invariant dynamics are assumed. Unfortunately it is difficult to write such a model, as the hydrodynamics of swimming are still a topic of current research, and the state x includes not only the robot's state but the infinite-dimensional state of the surrounding fluid. Even if we could capture the dynamic equations in the form (7), there do not exist techniques for designing controllers for general nonlinear systems of this form. More structure is needed before analytical design techniques can be brought to bear.

To simplify the equations of motion, we let x be simply the state of the robot fish. Model inaccuracies due to this simplification can then be treated as disturbances to be compensated for in closed-loop control. In this case, the state x can be written as the robot's configuration q and velocity \dot{q} . The configuration q is written $q = (g, b) \in Q = G \times B$, where $g \in G$ is the position and orientation of a frame fixed to the robot fish in a world inertial frame, called the "group" variables, and $b \in B$ is a set of "shape" variables describing the robot's internal degrees-of-freedom, such as the joint angles of pectoral or caudal fins (Figure 3). Typically the body frame configuration space G is $SE(3)$, the six-dimensional matrix group of positions and orientations of a rigid body in three-dimensional space

$$g = \begin{bmatrix} R & p \\ 0 & 1 \end{bmatrix}, \quad p \in \mathbb{R}^3, R \in \mathbb{R}^{3 \times 3}, RR^T = I, \det(R) = 1,$$

where p is the position of the reference point of the robot and R is a 3×3 rotation matrix describing the robot's orientation in an inertial frame. The shape space B may be considered infinite-dimensional if we assume fins are made of flexible material, but we will remain finite-dimensional.

Although the body-fixed frame configuration g can be globally represented as an element of the matrix Lie group $SE(3)$, and its velocity can be properly thought of as an element of the Lie algebra $\mathfrak{se}(3)$, it is beyond the scope of this paper to review the use of Lie groups and Lie algebras in motion planning and control (see, for example, the texts [91], [54], [103], or the papers [114], [80], [60], [88] for their use in the context of robot fish, or the papers [67], [30] in the context of autonomous underwater vehicles). For this reason, we will assume that g is written as an element of \mathbb{R}^6 (or \mathbb{R}^3 if the fish is restricted to a plane) using local coordinates, such as x-y-z-roll-pitch-yaw, with the velocity given by the time-derivatives of these coordinates. We will continue to refer to g as the group variables.

A.1 Model Abstractions

In the development of control algorithms for robot arms, the control signals are joint torques and the controlled outputs are joint angles. This is an abstraction, however; the control algorithm actually sends a control signal to a motor amplifier, which implements its own feedback loop to try to generate a motor current proportional to the signal. This "inner" feedback loop is usually assumed to be available and is not part of the robot control system design.

In the design of a controller for a robot fish, we need to choose an appropriate level of abstraction. Our interest is in controlling the motion of the body frame $g(t)$, and these variables are indirectly controlled by their state-dependent coupling to the directly actuated shape degrees-of-freedom $b(t)$. At the lowest level of abstraction, the control signals are generalized forces at the robot fish's joints. At a higher level of abstraction, we could assume the control inputs are the joint accelerations $\ddot{b}(t)$. This would then require an inner feedback loop to realize the commanded joint motions (see, for example, the caudal fin trajectory controller described in [101], [102]). At a still higher level of abstraction, we could assume the robot fish is a rigid body and the control inputs are forces and torques on the body. For example, in the case of a propeller-powered AUV, the shape variables (the propeller angles) are usually eliminated completely from the description of the configuration, and the propellers are treated simply as force and torque sources. Propeller thrust control, as a function of the current water flow and propeller speed, has been studied in [50], [9], [8]. Similarly, when the moving surfaces of a robot fish are "small" relative to the size of the fish, e.g., pectoral fin locomotion, we might follow the same approach, separating the control of the fish into an outer control loop for the group variables and an inner control loop for (possibly time-averaged) thrust control. A great deal of simulation and experimental work is currently aimed at characterizing the force generating capabilities of oscillating foils under different operating conditions. We do not review this work here, but see for example the review paper [109], work on pectoral fin controllers by Kato [55], [56], and the Nektar oscillating fin thruster [1].

An abstraction is useful when it both captures the essential dynamics of the system and simplifies the design of the control algorithm. The control of robot fish adds layers of abstraction complexity as compared to the control of a robot arm. For an arm, we assume control of joint torques, which are turned into joint motions in a state-dependent way. For a robot fish, we assume control of joint torques, which are turned into joint motions in a state-dependent way, producing forces on the body in a state-dependent way, resulting in motions $g(t)$ in a state-dependent way. To avoid restrictions on the subsequent discussion, we will assume the lowest-level abstraction, where the control inputs u are generalized forces acting on the shape coordinates.

A.2 Underactuation

Because our actuators directly control $b(t)$ but only indirectly control $g(t)$, a natural question is whether the configuration g is controllable at all. Since $\dim(B) < \dim(Q)$, the robot fish system is trivially *underactuated* — there are fewer control vari-

ables than configuration variables. Our primary interest is in controlling the g variables, however, so in this paper we will call the robot fish underactuated if $\dim(B) < \dim(G)$ — there are fewer controls than group configuration variables. An example of such a robot fish is the planar carangiform swimmer of Burdick *et al.* [79], [81], [87], [88], which has two controlled degrees-of-freedom of the caudal fin ($\dim(B) = 2$) but three degrees-of-freedom of the body (x-y-yaw) in the plane ($\dim(G) = 3$).

Most real fish are *redundant* in the sense that the number of controlled joints is greater than $\dim(G) = 6$, and the joints themselves have redundant muscular actuation [92]. A system with $\dim(B) \geq \dim(G)$ may still be underactuated, however, if the coupling from the control inputs to certain g directions is weak (i.e., if thrust cannot be generated in certain body directions). If the robot fish can generate a force-torque combination in an arbitrary direction in the g frame, for example an AUV with six or more fixed-direction propellers, we call it *fully actuated*.

Controllability of the g variables via direct actuation of the b variables is sometimes called *fiber controllability* [59], [33]. This terminology arises from the decomposition of the configuration space Q into a shape space B and a “fiber” attached to each $b \in B$, corresponding to the group variables g (see, e.g., [2], [104], [21]). Fiber controllability refers to the ability to control the g variables while not worrying about how the b variables evolve. In other words, we only care about controlling the evolution of the system on G , not on the full configuration space Q .

For a system of the form (7), let $q, \dot{q} \in \mathbb{R}^n$, written as column vectors, and write the state $x = (q^T, \dot{q}^T)^T \in M = \mathbb{R}^{2n}$ and the control $u \in \mathcal{U} \subset \mathbb{R}^m$, where \mathcal{U} contains a neighborhood of the origin in \mathbb{R}^m . Let $\mathcal{R}^V(x, \leq T)$ be the set of reachable states beginning from x in time $0 < t \leq T$ by feasible trajectories satisfying (7) and remaining in a neighborhood V of x on M . Then the system (7) is

- *(globally) controllable from x* if any state x_1 is reachable from x in finite time; that is, $x_1 \in \mathcal{R}^M(x, \leq T)$ for some $T > 0$;
- *small-time locally controllable (STLC) from x* if $\mathcal{R}^V(x, \leq T)$ contains a neighborhood of x for any $T > 0$ and any neighborhood V of x (this definition only makes sense from zero velocity states $x = (q^T, 0^T)^T$ for a dynamic system with forces as inputs);
- *configuration controllable from $x = (q^T, 0^T)^T$* if any state $(q_1^T, 0^T)^T$ is reachable from x in finite time;
- *small-time locally configuration controllable (STLCC) from $x = (q^T, 0^T)^T$* if $\mathcal{R}^V(x, \leq T)$ contains a neighborhood of q on Q for any $T > 0$ and any neighborhood V of x .

Fiber controllability concepts are a natural extension of these concepts, with the projection from the full state $(q^T, \dot{q}^T)^T$ down to the group state $(g^T, \dot{g}^T)^T$.

Controllability tests for underactuated systems of the form (7) are a topic of current research, but for the particular case of *control affine* nonlinear control systems, much is known. The class of systems which can be written in this form, perhaps after a feedback or coordinate transformation, is quite broad and includes some mathematical models of fish locomotion. Such

systems can be written

$$\dot{x} = X_0(x) + \sum_{i=1}^m X_i(x)u_i, \quad u \in \mathcal{U} \subset \mathbb{R}^m, \quad (8)$$

where $X_0(x)$ is a *drift vector field* describing the unforced natural motion of the system and the $X_i(x), i = 1 \dots m$, are *control vector fields* describing the action of the controls on the system. The *Lie bracket* of two vector fields X_i and X_j is a new vector field, written $[X_i, X_j]$, describing the second-order approximation of the motion achieved by following X_i for time ϵ , X_j for time ϵ , $-X_i$ for time ϵ , then $-X_j$ for time ϵ , as $\epsilon \rightarrow 0$:

$$x(4\epsilon) = x(0) + \epsilon^2[X_i, X_j](x(0)) + O(\epsilon^3),$$

where

$$[X_i, X_j](x) = \left(\frac{\partial X_2}{\partial x} X_1 - \frac{\partial X_1}{\partial x} X_2 \right) (x). \quad (9)$$

Lie brackets play a vital role in controllability tests for control affine nonlinear control systems. A system (8) is said to satisfy the *Lie algebra rank condition (LARC)* at x if the drift and control vector fields, along with their iterated Lie brackets, span the $2n$ -dimensional space of all possible motions at x (the tangent space $T_x M$). In other words, local “wiggling” maneuvers can break the apparent motion limitations due to underactuation. As an example, consider a robot fish with a configuration $q = (x, y, \theta)^T$, describing the position and orientation of the body in a horizontal plane. For simplicity, we will treat it as a kinematic system with velocities as inputs instead of forces (see, for example, [67], [68], where this assumption is used for AUV’s). This robot fish is drift-free ($X_0 = 0$) with two control vector fields: swim forward (surge), $X_1(q) = (\cos \theta, \sin \theta, 0)^T$, and turn in place (yaw), $X_2(q) = (0, 0, 1)^T$. Then the control system is written

$$\dot{q} = X_1(q)u_1 + X_2(q)u_2, \quad u_1, u_2 \in \mathbb{R},$$

and we calculate the Lie bracket

$$\begin{aligned} [X_1, X_2] &= \frac{\partial X_2}{\partial x} X_1 - \frac{\partial X_1}{\partial x} X_2 & (10) \\ &= \mathbf{0} \begin{pmatrix} \cos \theta \\ \sin \theta \\ 0 \end{pmatrix} - \begin{pmatrix} 0 & 0 & -\sin \theta \\ 0 & 0 & \cos \theta \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} \sin \theta \\ -\cos \theta \\ 0 \end{pmatrix}. \end{aligned}$$

This is a motion in the slip direction, or “parallel-parking” (Figure 4). The three vector fields X_1, X_2 , and $[X_1, X_2]$ span \mathbb{R}^3 at any q , so motion is locally possible in any direction despite the underactuation. Some systems require iterated Lie brackets, for example of the form $[X_1, [X_1, [X_1, X_2]]]$, to establish the LARC.

The LARC is the kernel of most tests for local controllability. If the LARC is satisfied at x , the system is *small-time locally accessible* at x , meaning $\mathcal{R}^V(x, \leq T)$ contains an open




Fig. 4. Generating motion in a Lie bracket direction.

set on M for any neighborhood V and any $T > 0$. If the system is drift-free, the LARC implies STLCC, under the assumption that \mathcal{U} contains a neighborhood of the origin of \mathbb{R}^m . For dynamic systems with drift ($X_0 \neq 0$), STLCC at zero-velocity states $x = (q^T, 0^T)^T$ can be established by the LARC and auxiliary conditions that certain “bad” Lie brackets can be expressed as linear combinations of “good” Lie brackets of lower degree [107]. The weaker condition of STLCC, which only requires that the locally reachable set contain a neighborhood of the initial configuration, not necessarily the velocity, can be established for a class of mechanical control systems by Lie algebraic tests described in [69], [70], and the corresponding fiber controllability tests can be found in [33].

The LARC is the nonlinear analog of the Kalman rank condition for linear systems. In general, the linearized dynamics of an underactuated robot fish will be uncontrollable. Therefore, controllability for an underactuated system is inherently a *nonlinear* phenomenon. This has important consequences in the design of feedback controllers and trajectory planners.

B. Station Keeping

Station keeping is the problem of maintaining a desired equilibrium state in the presence of disturbances. For a robot fish control system of the form

$$\dot{x} = f(x, u), \quad x \in M, \quad u \in \mathbb{R}^m,$$

let x_0 be the desired equilibrium state, such that $f(x_0, 0) = 0$. A necessary condition for the existence of a continuous state feedback law $u(x)$ to asymptotically stabilize x_0 is that

$$f(\mathcal{B}, \mathbb{R}^m) = \{f(x, u) | x \in \mathcal{B}, u \in \mathbb{R}^m\}$$

must contain a neighborhood of the origin in M for any neighborhood \mathcal{B} of x_0 on M [23], [32]. If the system is underactuated, then this condition will not be satisfied. This observation has led to the development of feedback controllers for underactuated systems which are time-varying, i.e., of the form $u(x, t)$, or discontinuous, i.e., of the form $u_{i(x)}(x)$, where $i : M \rightarrow \{1, \dots, p\}$ is an index function which chooses among a set of p controllers based on the current state. One approach to implementing a time-varying controller is to plan a trajectory to the goal state, execute it, and iterate the process, ensuring that the final error decreases after each iterate. Feedback is only incorporated between iterates. Another approach is to continually re-plan and immediately begin to implement the new trajectory, before the previous one is finished. This is a type of model predictive control. Trajectory planning is described in the next section.

Any disturbance that pushes the robot fish in a non-actuated direction can only be recovered from by locally executing a maneuver (e.g., the parallel-parking slip maneuver of Section III-A.2), which may be time-consuming. For this reason, it may be preferable to design a robot fish which can thrust in any direction when station holding is a primary function. Fish employing MPF swimming modes can generate a greater range of thrust directions than those employing BCF swimming modes, which may be why many fish operating in littoral zones rely on median and pectoral fins.

C. Trajectory Planning and Tracking

A primary control goal is to be able to drive the robot fish from one state to another, possibly in the presence of control constraints and obstacles. **As with station keeping, underactuation of the robot fish results in considerable challenges in control design. The linearization of the fish dynamics is not generally controllable, so linear feedback control design is inapplicable.** Discontinuous or time-varying feedback controllers may be designed to drive the system to a goal state, but pure feedback controllers are likely to result in inefficient trajectories. If it is not necessary to completely control the group variables g , it may be possible to choose a subset of these variables as the controlled “outputs” and perform an input-output feedback linearization. In this case, care must be taken to ensure that the uncontrolled state variables, the “zero dynamics,” are stable.

An alternative approach to control of state-to-state motions for an underactuated system has been called a *two degree-of-freedom* approach [73], [111], [44]. In this approach, a nominal trajectory for the system is planned in advance, and a feedback controller is used to track the trajectory. The feedforward control derived by the trajectory planner accounts for the nominal nonlinear dynamics of the system, while the feedback controller accounts for disturbances and model inaccuracies. This feedback controller may be a linear controller, since the linearized system is likely to be controllable about a trajectory, despite not being controllable at an equilibrium.

The goal of the trajectory planner is to find a time- or energy-efficient state-to-state trajectory for the robot fish, and to find that trajectory in a computationally efficient manner, possibly allowing real-time implementation. The first issue can be addressed by numerical optimal control methods such as nonlinear optimization or dynamic programming [66], [25], [71], [106]. Nonlinear optimization can be used to solve for the optimal control parameters in a finite parameterization of the input histories. The search for the optimal parameters in the design space is usually aided by gradient information, assuming the objective function and control and state constraints are sufficiently smooth with respect to the control parameters. The result is a locally optimal control in the design space. This approach was used by Saimek and Li [101], [102] to find energy-efficient motion primitives for a caudal fin swimmer, and by MacIver [76] to find “effort” minimizing motions for a fully-actuated ellipsoidal body in water, producing motions mimicking the predatory behavior of a gymnotid knife fish.

Numerical optimal control methods address the goal of generating time-optimal or energy-efficient motion, but they tend to be computationally intensive. On the other hand, randomized

techniques from the robot motion planning literature [57], [65] attempt to quickly find satisficing trajectories without concern for optimality. These techniques are particularly applicable to systems operating in cluttered environments.

Another approach to motion planning for robot fish is based on *motion libraries*. This approach is based on the observation that many biological systems employ a small number of stereotypical maneuvers, motion primitives, or gaits, and these primitives are efficient in some way. For example, fish locomotion is characterized by highly stereotypical oscillatory or undulatory motions of the fins and body. A motion library consists of a set of primitive maneuvers which are concatenable, so that a trajectory plan consists of a specification of a sequence of maneuvers and their switching conditions. One formalization of this idea is the *maneuver automaton* of Frazzoli *et al.* [41], which consists of a set of motion primitives in a graph structure indicating how they can be concatenated. Motion libraries have also been used to generate natural motions for animated characters [62]. Motion libraries are an instantiation of the general idea of *hierarchical motion planning*, where the task of planning the motion of a complex mechanical system is broken down into simpler planning tasks for reduced-complexity approximations of the full system [98], [100], [30], [24], [37], [52], [78]. Example motion libraries for robot fish locomotion have been generated by Ayers *et al.* [7], Wilbur *et al.* [120], and Ostrowski and McIsaac [98], [83] for anguilliform locomotion; Burdick *et al.* [81], [87], [88] and Saimek and Li [102], [101] for carangiform locomotion; and Kato [55], [56] for pectoral fin locomotion.

The primitives of the motion library may be in the form of feedback controllers or open-loop input profiles. The primitives are generated off-line, allowing experimental derivation of the primitives, or the use of time-consuming numerical optimal control methods to find the primitives. To limit the number of primitives stored, it is useful to take advantage of symmetries in the dynamics (7) to create parameterized primitives. For example, the primitive `swim-forward(d)`, where d is the total distance, is more useful than `swim-forward-2meters`. Since we expect the dynamics to be invariant to displacements of the fish in a horizontal plane, such a parameterized primitive is possible.

Perhaps more interestingly, systems satisfying the potential flow assumption of (6) satisfy a time-scaling property: any feasible trajectory $q(t)$ implies a family of feasible trajectories of the form $q(\lambda t)$ [51], [102], [101]. A choice $\lambda > 1$ speeds up the trajectory, while $0 < \lambda < 1$ slows down the trajectory. Therefore, our primitive above could be modified to have two parameters, `swim-forward(d, τ)`, where τ is the time of motion. The required input magnitudes scale with λ^2 . If the primitive was derived to be optimal for $\lambda = 1$, it may be far from optimal for λ values significantly different from 1. This might suggest multiple `swim-forward` primitives corresponding to qualitatively different “gaits” of the fish. Velocity-dependent transitions between gaits can be triggered by efficiency considerations or speed range limitations for each gait, much as in real fish and other biological systems [36], [61], [40], [5].

For some underactuated systems of the form (6), certain trajectories $q(t)$ permit more general time-scalings than simple

uniform time-scalings $q(\lambda t)$. A trajectory $q(t)$ is called a *kinematic motion* for the dynamic system if there exists a control u satisfying (6) for the trajectory $q(s(t))$ for any twice-differentiable time-scaling function $s(t) \in \mathbb{R}$ [75], [30], [29]. Any velocity vector field $X(q)$ whose integral curves are kinematic motions is called a *decoupling vector field*, meaning that motion planning along such a vector field can be decoupled into a kinematic problem (deciding how far to follow the vector field) followed by a dynamic problem (deciding the time-scaling, or how fast to follow the vector field, subject to actuator limits). These vector fields can be thought of as “forced” trim trajectories or relative equilibria, and in the case of vehicles often correspond to body-fixed velocities. The advantage of decoupling vector fields in motion planning is that planning a robot trajectory, possibly amidst obstacles, can be reduced from a dynamic problem to a kinematic problem (plus time-scaling), reducing the dimension of the search space by a factor of two and resulting in considerable computational savings. This approach was applied to an underactuated AUV model in [30]. Decoupling vector fields can be used as motion library primitives, suggested directly by the dynamics of the system.

More generally, if the system can be written in the control affine form (8), it is well known that time-periodic inputs can be used to generate motions in directions corresponding to the Lie brackets of the system vector fields. For example, plugging in sinusoidal inputs and integrating the equations of motion, we find that the net motion at the end of a cycle can be expressed as a series expansion in terms of iterated Lie brackets. The contributions of each of the Lie bracket terms to the total motion depends on the magnitude, frequency, and phase of the sinusoidal inputs. The side-slip motion of the kinematic robot fish of Figure 4, for example, can be approximately generated by 90° out-of-phase sinusoidal controls, instead of the alternating on-off controls depicted in the figure.

Averaging formulas have been derived by many to describe the average effect of sinusoidal inputs over a cycle [48], [68], [67], [10], [97], [93], [28], [26], [27], [86], [113], [114], [112], [88], [21]. This approach is appealing because fish locomotion often involves oscillatory motions of the fins and body. These formulas also provide a way to analytically find controls which approximately generate motions in “surprising” directions; for example, a carangiform robot fish can be made to approximately rotate in place using only the caudal fin [87]. A drawback of the averaging approach is that the formulas often only apply when the robot is traveling at low speeds. Also, motions in Lie bracket directions tend to be slow; consider, for example, how tedious it is to parallel park a car. Thus, even though such motions are possible, they may not be time- or energy-efficient, and should be properly penalized in any trajectory planning scheme.

Once a nominal trajectory has been generated, the problem is to track it. The nominal model of the dynamics provides the feedforward controls, and feedback control is used to compensate for perturbations and modeling error. One approach is to linearize the dynamics about the trajectory and design a linear controller based on the error coordinates. Another approach, proposed in [88], [113], [112], uses averaging formulas based on the system’s Lie brackets to construct controllers that use discrete-time state feedback (at the end of each sinusoidal con-

trol cycle) to choose controls to drive the system back to the planned trajectory (or equilibrium state).

Because of the complexity of the hydrodynamics of swimming, one possibility is to give up on simple, low-order descriptions of the dynamic equations, and instead to use rule-based models or empirical, data-driven models in trajectory planning for robot fish [6]. Barrett [13] used extensive experiments with the RoboTuna to find a motion pattern optimizing the thrust power ratio, defined as the thrust power divided by the power input to undulate the body. The motion pattern is described by seven parameters, such as the forward velocity, the Strouhal number, the amplitude of the caudal fin motion, etc. To efficiently guide the search through the seven-parameter space without using gradient information (because of possible discontinuities in hydrodynamic regimes, e.g., the transition from laminar to turbulent flow), a genetic algorithm (GA) was used. The GA was specially designed for this problem to minimize the number of expensive experiments needed for convergence. Grzeszczuk and Terzopoulos [46] used another non-smooth global optimization method, simulated annealing, to discover muscle activation patterns resulting in efficient, life-like swimming gaits for simulated fish.

Galls and Redionitis [43] use a CFD simulation of carangiform swimming to generate test cases for training a relatively low order neural network controller. One difficulty with this approach is that, as the number of control inputs and the dependence on control history grow, the number of test cases becomes prohibitive.

Kato [55], [56] deals with the complexity of swimming hydrodynamics by generating fuzzy rule-based control laws for point-to-point motions of a robot fish with two rigid pectoral fins, each with two or three degrees-of-freedom. For the control law currently in operation, fuzzy rules based on error coordinates are used to choose the parameters of sinusoidal motions of the pectoral fins. These control laws have been implemented to drive a robot fish both in a horizontal plane and in three dimensions.

IV. CONCLUSION

Achieving fish-like capabilities in maneuverability and efficiency in an AUV will require the tight integration of design and control development. This effort will benefit from new insight from biology, tractable and accurate hydrodynamics models, new actuation and sensing technology, and advances in non-linear control theory and our understanding of the neural mechanisms of control. This paper has reviewed lessons learned so far on the design of swimming machines and the control problems that must be solved to make full use of their capabilities.

REFERENCES

- [1] <http://www.nektonresearch.com>.
- [2] R. Abraham, J. E. Marsden, and T. S. Ratiu. *Manifolds, Tensor Analysis, and Applications*. Springer-Verlag, 1988.
- [3] R. M. Alexander. *Elastic mechanisms in animal movement*. Cambridge University Press, 1988.
- [4] J. M. Anderson and N. K. Chhabra. Maneuvering and stability performance of a robotic tuna. *Integrative and Comparative Biology*, 42(1):118–126, 2002.
- [5] V. I. Arreola and M. W. Westneat. Mechanics of propulsion by multiple fins: Kinematics of aquatic locomotion in the burrfish (*Chilomycterus schoepfi*). *Biological Sciences*, 263:1689–1696, 1996.
- [6] C. G. Atkeson and S. Schaal. Memory-based neural networks for robot learning. *Neurocomputing*, 9(3):243–269, 1995.
- [7] J. Ayers, C. Wilbur, and C. Olcott. Lamprey robots. In T. Wu and N. Kato, editors, *Proceedings of the International Symposium on Aqua Biomechanisms*. Tokai University, 2000.
- [8] R. Bachmayer and L. L. Whitcomb. Adaptive parameter identification of an accurate nonlinear dynamical model for marine thrusters. in press.
- [9] R. Bachmayer, L. L. Whitcomb, and M. A. Grosenbaugh. An accurate four-quadrant nonlinear dynamical model for marine thrusters: Theory and experimental validation. *IEEE Journal of Oceanographic Engineering*, 25(1):146–159, Jan. 2000.
- [10] J. Baillieul and B. Lehman. Open-loop control using oscillatory inputs. In *CRC Control Handbook*, pages 967–980. CRC Press, Boca Raton, FL, 1996.
- [11] P. Bandyopadhyay, J. Castano, J. Rice, R. Phillips, W. Nedderman, and W. Macy. Low speed maneuvering hydrodynamics of fish and small underwater vehicles. *ASME J. Fluids Eng.*, 119:136–144, 1997.
- [12] P. R. Bandyopadhyay. Maneuvering dynamics of fish and small underwater vehicles. *Integrative and Comparative Biology*, 42(1):102–117, 2002.
- [13] D. Barrett. Optimization of swimming locomotion by genetic algorithm. In J. Ayers, J. L. Davis, and A. Rudolph, editors, *Neurotechnology for Biomimetic Robots*. MIT Press, 2002.
- [14] D. S. Barrett, M. S. Triantafyllou, D. K. P. Yue, M. A. Grosenbaugh, and M. J. Wolfgang. Drag reduction in fish-like locomotion. *J. Fluid Mech.*, 392:183–212, 1999.
- [15] I. Bartol, M. Gharib, D. Weihs, P. Webb, J. Hove, and M. Gordon. Hydrodynamic stability of swimming in ostaciid fishes: role of the carapace in the smooth trunkfish *Lactophrys triqueter* (teleostei: Ostraciidae). *J. Exp. Biol.*, 206:725–744, 2003.
- [16] R. W. Blake. The energetics of hovering in the mandarin fish. *J. Exp. Biol.*, 82:25–33, 1979.
- [17] R. W. Blake. The mechanics of labriform locomotion I. Labriform locomotion in the angelfish (*pterophyllum eimekei*): an analysis of the power stroke. *J. Exp. Biol.*, 82:255–271, 1979.
- [18] R. W. Blake. The mechanics of labriform locomotion II. An analysis of the recovery stroke and the overall fin-beat cycle propulsive efficiency in the angelfish. *J. Exp. Biol.*, 85:337–342, 1980.
- [19] R. W. Blake. Swimming in the electric-eels and knifefishes. *Can. J. Zool.*, 61:1432–1441, 1983.
- [20] A. R. Blight. The muscular control of vertebrate swimming movements. *Biol. Rev.*, 52:181–218, 1977.
- [21] A. M. Bloch. *Nonholonomic Mechanics and Control*. Springer, New York, 2003.
- [22] M. G. Borgen, G. N. Washington, and G. L. Kinzel. Design and evolution of a piezoelectrically actuated miniature swimming vehicle. *IEEE/ASME Transactions on Mechatronics*, 8(1):66–74, Mar. 2003.
- [23] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. S. Millman, and H. J. Sussmann, editors, *Differential Geometric Control Theory*. Birkhauser, 1983.
- [24] R. W. Brockett. Hybrid models for motion control systems. In *Essays in Control: Perspectives in the Theory and its Applications*, pages 29–53. Birkhauser, 1993.
- [25] A. E. Bryson. *Dynamic Optimization*. Addison-Wesley, 1998.
- [26] F. Bullo. Series expansions for the evolution of mechanical control systems. *SIAM Journal on Control and Optimization*, 40(1):166–190, 2001.
- [27] F. Bullo. Averaging and vibrational control of mechanical systems. *SIAM Journal on Control and Optimization*, 41:542–562, 2002.
- [28] F. Bullo, N. E. Leonard, and A. D. Lewis. Controllability and motion algorithms for underactuated Lagrangian systems on Lie groups. *IEEE Transactions on Automatic Control*, 45(8):1437–1454, 2000.
- [29] F. Bullo, A. D. Lewis, and K. M. Lynch. Controllable kinematic reductions for mechanical systems: Concepts, computational tools, and examples. In *2002 International Symposium on the Mathematical Theory of Networks and Systems*, Aug. 2002.
- [30] F. Bullo and K. M. Lynch. Kinematic controllability for decoupled trajectory planning of underactuated mechanical systems. *IEEE Transactions on Robotics and Automation*, 17(4):402–412, Aug. 2001.
- [31] S. Coombs, J. J. Finneran, and R. A. Conley. Hydrodynamic image formation by the peripheral lateral line system of the Lake Michigan mottled sculpin, *Cottus bairdi*. *Phil. Trans. R. Soc. Lond. B*, 355:1111–1114, 2000.
- [32] J.-M. Coron. A necessary condition for feedback stabilization. *Systems and Control Letters*, 14:227–232, 1990.
- [33] J. Cortes, S. Martinez, J. P. Ostrowski, and H. Zhang. Simple mechanical control systems with constraints and symmetry. *SIAM Journal on Control and Optimization*, 41(3):851–874, 2002.

- [34] M. H. Dickinson. Unsteady mechanisms of force generation in aquatic and aerial locomotion. *Amer. Zool.*, 36:537–554, 1996.
- [35] S. Dijkgraaf. The functioning and significance of the lateral-line organs. *Biol. Rev.*, 38:51–105, 1962.
- [36] E. G. Drucker. The use of gait transition speed in comparative studies of fish locomotion. *American Zoologist*, 36:555–566, 1996.
- [37] M. Egerstedt. Motion description languages for multi-modal control in robotics. In A. Bicchi, H. I. Christensen, and D. Prattichizzo, editors, *Control Problems in Robotics*. Springer, 2003.
- [38] Z. Fan, J. Chen, J. Zou, J. Li, C. Liu, and F. Delcomyn. Development of artificial lateral-line flow sensors. In *Solid-State Sensor, Actuator and Microsystems Workshop*, Hilton Head Island, South Carolina, 2002.
- [39] F. Fish. Balancing requirements for stability and maneuverability in cetaceans. *Integrative and Comparative Biology*, 42(1):85–93, 2002.
- [40] F. E. Fish. Transitions from drag-based to lift-based propulsion in mammalian swimming. *American Zoologist*, 36:628–641, 1996.
- [41] E. Frazzoli, M. A. Dahleh, and E. Feron. Real-time motion planning for agile autonomous vehicles. *AIAA Journal of Guidance, Control, and Dynamics*, 25(1):116–129, 2002.
- [42] Y. C. Fung. *An Introduction to the Theory of Aeroelasticity*. John Wiley and Sons, New York, 1955.
- [43] S. F. Galls and O. K. Rediniotis. Simulation of fish locomotion and control. In *AIAA Paper No. 2000-0296*. 38th Aerospace Science Meeting and Exhibit, Jan. 2000.
- [44] A. Giusto and F. Paganini. Robust synthesis of feedforward compensators. *IEEE Transactions on Automatic Control*, 44(8):1578–82, 1999.
- [45] M. Gordon, J. Hove, P. Webb, and D. Weihs. Boxfishes as unusually well-controlled autonomous underwater vehicles. *Physiological and Biochemical Zoology*, 74(6):663–671, 2000.
- [46] R. Grzeszczuk and D. Terzopoulos. Automated learning of muscle-actuated locomotion through control abstraction. In *SIGGRAPH 95*, pages 63–70, 1995.
- [47] S. Guo, T. Fukuda, and K. Asaka. A new type of fish-like underwater microrobot. *IEEE/ASME Transactions on Mechatronics*, 8(1):136–141, Mar. 2003.
- [48] L. Gurvits. Averaging approach to nonholonomic motion planning. In *IEEE International Conference on Robotics and Automation*, pages 2541–2546, 1992.
- [49] K. A. Harper, M. D. Berkemeier, and S. Grace. Modeling the dynamics of spring-driven oscillating-foil propulsion. *IEEE Journal of Oceanic Engineering*, 23(3):285–296, July 1998.
- [50] A. J. Healey, S. M. Rock, S. Cody, D. Miles, and J. P. Brown. Toward an improved understanding of thruster dynamics for underwater vehicles. *IEEE Journal of Oceanographic Engineering*, 20:354–361, Oct. 1995.
- [51] J. M. Hollerbach. Dynamic scaling of manipulator trajectories. *ASME Journal of Dynamic Systems, Measurement, and Control*, 106:102–106, 1984.
- [52] D. Hristu and S. Andersson. Directed graphs and motion description languages for robot navigation and control. In *IEEE International Conference on Robotics and Automation*, 2002.
- [53] J. John H. Long, M. E. Hale, M. J. McHenry, and M. W. Westneat. Functions of fish skin: flexural stiffness and steady swimming of longnose gar *Lepisosteus Osseus*. *J. Exp. Biol.*, 199:2139–2151, 1996.
- [54] V. Jurdjevic. *Geometric Control Theory*. Cambridge University Press, 1997.
- [55] N. Kato. Control performance in the horizontal plane of a fish robot with mechanical pectoral fins. *IEEE Journal of Oceanic Engineering*, 25(1):121–129, Jan. 2000.
- [56] N. Kato. Pectoral fin controllers. In J. Ayers, J. L. Davis, and A. Rudolph, editors, *Neurotechnology for Biomimetic Robots*. MIT Press, 2002.
- [57] L. Kavraki, P. Švestka, J.-C. Latombe, and M. Overmars. Probabilistic roadmaps for fast path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12:566–580, 1996.
- [58] S. D. Kelly, R. J. Mason, C. T. Anhalt, R. M. Murray, and J. W. Burdick. Modelling and experimental investigation of carangiform locomotion for control. In *American Control Conference*, pages 1271–1276, Philadelphia, PA, 1998.
- [59] S. D. Kelly and R. M. Murray. Geometric phases and robotic locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995.
- [60] S. D. Kelly and R. M. Murray. Modelling efficient pisciform swimming for control. *International Journal on Nonlinear and Robust Control*, 10(4):217–241, Apr. 2000.
- [61] K. E. Korsmeyer, J. F. Steffensen, and J. Herskin. Energetics of median and paired fin swimming, body and caudal fin swimming, and gait transition in parrotfish (*Scarus schlegelii*) and triggerfish (*Rhinecanthus aculeatus*). *Journal of Experimental Biology*, 205:1253–1263, 2002.
- [62] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. *ACM Transactions on Graphics*, 21(3), July 2002.
- [63] G. V. Lauder and E. G. Drucker. Morphology and experimental hydrodynamics of piscine control surfaces. Preprint.
- [64] G. V. Lauder and E. G. Drucker. Forces, fishes, and fluids: Hydrodynamic mechanisms of aquatic locomotion. *News in Physiological Sciences*, 17:235–240, 2002.
- [65] S. M. LaValle and J. J. Kuffner. Rapidly-exploring random trees: Progress and prospects. In B. R. Donald, K. M. Lynch, and D. Rus, editors, *Algorithmic and Computational Robotics: New Directions*. A. K. Peters, Natick, MA, 2001.
- [66] E. B. Lee and L. Markus. *Foundations of Optimal Control Theory*. Krieger Publishing Company, 1986.
- [67] N. E. Leonard. Control synthesis and adaptation for an underactuated autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 20(3):211–220, July 1995.
- [68] N. E. Leonard and P. S. Krishnaprasad. Motion control of drift-free, left-invariant systems on Lie groups. *IEEE Transactions on Automatic Control*, 40(9):1539–1554, Sept. 1995.
- [69] A. D. Lewis and R. M. Murray. Configuration controllability of simple mechanical control systems. *SIAM Journal on Control and Optimization*, 35(3):766–790, May 1997.
- [70] A. D. Lewis and R. M. Murray. Configuration controllability of simple mechanical control systems. *SIAM Review*, 41(3):555–574, 1999.
- [71] F. L. Lewis and V. L. Syrmos. *Optimal Control*. John Wiley and Sons, Inc., 1995.
- [72] P. Y. Li and S. Saimek. Modeling and estimation of hydrodynamic potentials. In *IEEE International Conference on Decision and Control*, 1999.
- [73] D. J. N. Limebeer, E. M. Kasenally, and J. D. Perkins. On the design of robust two degrees of freedom controllers. *Automatica*, 29(1), 1993.
- [74] J. H. Long and K. S. Nipper. The importance of body stiffness in undulatory propulsion. *American Zoologist*, 36:678–694, 1996.
- [75] K. M. Lynch, N. Shiroma, H. Arai, and K. Tanie. Collision-free trajectory planning for a 3-DOF robot with a passive joint. *International Journal of Robotics Research*, 19(12):1171–1184, Dec. 2000.
- [76] M. A. MacIver. Neuromechanical design and active sensory systems in animals. Preprint.
- [77] M. A. MacIver, N. M. Sharabash, and M. E. Nelson. Prey-capture behavior in gymnotid electric fish: Motion analysis and effects of water conductivity. *J. Exp. Biol.*, 204(3):543–557, 2001.
- [78] V. Manikonda, P. S. Krishnaprasad, and J. Hendler. Languages, behaviors, hybrid architectures and motion control. In J. Baillieul and J. C. Willems, editors, *Mathematical Control Theory*. Springer-Verlag, 1998.
- [79] R. Mason and J. W. Burdick. Construction and modelling of a carangiform robotic fish. In *International Symposium on Experimental Robotics*, 1999.
- [80] R. Mason and J. W. Burdick. Propulsion and control of deformable bodies in an ideal fluid. In *IEEE International Conference on Robotics and Automation*, 1999.
- [81] R. Mason and J. W. Burdick. Experiments in carangiform robotic fish locomotion. In *IEEE International Conference on Robotics and Automation*, 2000.
- [82] M. M. McHenry, C. A. Pell, and J. H. Long, Jr. Mechanical control of swimming speed: stiffness and axial wave form in undulating fish models. *J. Exp. Biol.*, 198:2293–2305, 1995.
- [83] K. P. McIsaac and J. P. Ostrowski. Experiments in closed-loop control for an underwater eel-like robot. In *IEEE International Conference on Robotics and Automation*, pages 750–755, 2002.
- [84] R. Mittal. Computational modeling in bio-hydrodynamics. Preprint.
- [85] J. C. Montgomery, C. F. Baker, and A. G. Carton. The lateral line can mediate rheotaxis in fish. *Nature*, 389:960–963, Oct. 1997.
- [86] K. A. Morgansen. *Temporal patterns in learning and control*. PhD thesis, Harvard University, 1999.
- [87] K. A. Morgansen, V. Duindam, R. J. Mason, J. W. Burdick, and R. M. Murray. Nonlinear control methods for planar carangiform robot fish locomotion. In *IEEE International Conference on Robotics and Automation*, 2001.
- [88] K. A. Morgansen, P. A. Vela, and J. W. Burdick. Trajectory stabilization for a planar carangiform robot fish. In *IEEE International Conference on Robotics and Automation*, 2002.
- [89] P. Moyle and J. J. Cech, Jr. *Fishes: An Introduction to Ichthyology*. Prentice Hall, Upper Saddle River, New Jersey, 2000.
- [90] U. K. Müller, B. L. E. V. den Heuvel, E. J. Stamhuis, and J. J. Videler. Fish foot prints: Morphology and energetics of the wake behind continuously swimming mullet (*Chelon Labrosus Risso*). *J. Exp. Biol.*, 200:2893–2906, 1997.
- [91] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

- [92] F. A. Mussa-Ivaldi and S. A. Solla. Neural primitives for motion control. Preprint.
- [93] Y. Nakamura, T. Suzuki, and M. Koinuma. Nonlinear behavior and control of a nonholonomic free-joint manipulator. *IEEE Transactions on Robotics and Automation*, 13(6):853–862, 1997.
- [94] M. Nakashima and K. Ono. Development of a two-joint dolphin robot. In J. Ayers, J. L. Davis, and A. Rudolph, editors, *Neurotechnology for Biomimetic Robots*. MIT Press, 2002.
- [95] J. G. New, L. A. Fewkes, and A. N. Khan. Strike feeding behavior in the muskellunge, *Esox masquinongy*: Relative contributions of lateral line and visual sensory systems. *J. Exp. Biol.*, 204:1207–1221, 2001.
- [96] O. O. K. Rediniotis and D. C. Lagoudas. Shape memory alloy actuators as locomotor muscles. In *Workshop on Fixed, Flapping and Rotating Wing Micro-Air Vehicles*. University of Notre Dame, South Bend, Indiana, June 2000.
- [97] J. P. Ostrowski and J. W. Burdick. The geometric mechanics of undulatory robotic locomotion. *International Journal of Robotics Research*, 17(7):683–701, July 1998.
- [98] J. P. Ostrowski and K. Mclsaac. A framework for steering dynamic robotic locomotion systems. In B. R. Donald, K. M. Lynch, and D. Rus, editors, *Algorithmic and Computational Robotics: New Directions*. A. K. Peters, Natick, MA, 2001.
- [99] D. A. Pabst. Springs in swimming animals. *American Zoologist*, 36:723–735, 1996.
- [100] G. J. Pappas and S. Simic. Consistent abstractions of affine control systems. *IEEE Transactions on Automatic Control*, 47(5):745–756, May 2002.
- [101] S. Saimek and P. Y. Li. Longitudinal motion planning and control of a swimming machine. Submitted to *International Journal of Robotics Research*.
- [102] S. Saimek and P. Y. Li. Motion planning and control of a swimming machine. In *American Control Conference*, 2001.
- [103] S. S. Sastry. *Nonlinear Systems: Analysis, Stability, and Control*. Springer-Verlag, New York, 1999.
- [104] B. Schutz. *Geometrical methods of mathematical physics*. Cambridge University Press, 1980.
- [105] M. Sfakiotakis, D. M. Lane, and J. B. C. Davies. Review of fish swimming modes for aquatic locomotion. *IEEE J. Oceanic Eng.*, 24(2):237–252, 1999.
- [106] R. F. Stengel. *Optimal control and estimation*. Dover, New York, 1994.
- [107] H. J. Sussmann. A general theorem on local controllability. *SIAM Journal on Control and Optimization*, 25(1):158–194, Jan. 1987.
- [108] M. Triantafyllou, G. Triantafyllou, and R. Gopalkrishnan. Wake mechanics for thrust generation in oscillating foils. *Phys. Fluids A*, 3:2835–2837, 1991.
- [109] M. S. Triantafyllou, A. H. Techet, and F. S. Hover. Review of experimental work in biomimetic foils. Preprint.
- [110] M. S. Triantafyllou and G. S. Triantafyllou. An efficient swimming machine. *Scientific American*, pages 64–70, Mar. 1995.
- [111] M. J. van Nieuwstadt and R. M. Murray. Real time trajectory generation for differentially flat systems. *International Journal on Robust and Nonlinear Control*, 8(11):995–1020, 1998.
- [112] P. Vela and J. W. Burdick. Control of biomimetic locomotion via averaging theory. In *IEEE International Conference on Robotics and Automation*, 2003.
- [113] P. Vela, K. A. Morgansen, and J. W. Burdick. Second order averaging methods and oscillatory feedback control of underactuated mechanical systems, 2002. Submitted to *SIAM Journal on Control and Optimization*.
- [114] P. A. Vela, K. A. Morgansen, and J. W. Burdick. Underwater locomotion from oscillatory shape deformations. In *IEEE International Conference on Decision and Control*, 2002.
- [115] J. A. Walker and M. W. Westneat. Mechanical performance of aquatic rowing and flying. *Proceedings of the Royal Society of London B*, 267:1875–1881, 2000.
- [116] P. W. Webb. Control of posture, depth and swimming trajectories for fishes. *Integrative and Comparative Biology*, 42(1):94–101, 2002.
- [117] P. W. Webb and D. Weihs. Hydrostatic stability of fish with swimbladders: Not all fish are unstable. *Can. J. Zool.*, 72:1149–1154, 1994.
- [118] D. Weihs. Stability versus maneuverability in aquatic locomotion. *Integrative and Comparative Biology*, 42(1):127–134, 2002.
- [119] M. W. Westneat, D. H. Thorsen, J. A. Walker, and M. E. Hale. Structure, function and neural control of pectoral fins in fishes. Preprint.
- [120] C. Wilbur, W. Vorus, Y. Cao, and S. N. Currie. A lamprey-based undulatory vehicle. In J. Ayers, J. L. Davis, and A. Rudolph, editors, *Neurotechnology for Biomimetic Robots*. MIT Press, 2002.