

Multi-objective Optimization of Artificial Swimmers

Siddhartha Verma*, Panagiotis Hadjidoukas*, Philipp Wirth* and Petros Koumoutsakos*^{†‡}

*Computational Science and Engineering Laboratory, ETH Zürich, Clausiusstrasse 33, CH-8092, Switzerland

[†]Radcliffe Institute of Advanced Study, Harvard University, MA, United States of America

[‡]Wallace Visiting Professor, Massachusetts Institute of Technology, MA, United States of America

Email for Correspondence: petros@ethz.ch

Abstract—A fundamental understanding of how various biological traits and features provide organisms with a competitive advantage can help us improve the design of several mechanical systems. Numerical optimization can be invaluable for this purpose, by allowing us to scrutinize the evolution of specific biological adaptations. Importantly, the use of numerical optimization can help us overcome limiting constraints that restrict the evolutionary capability of biological species. Thus, we couple high-fidelity simulations of self-propelled swimmers with evolutionary optimization algorithms, to examine peculiar swimming patterns observed in a number of fish species. More specifically, we investigate the intermittent form of locomotion referred to as ‘burst-and-coast’ swimming, which involves a few quick flicks of the fish’s tail followed by a prolonged unpowered glide. This mode of swimming is believed to confer energetic benefits, in addition to several other advantages. We discover a range of intermittent-swimming patterns, the most efficient of which resembles the swimming-behaviour observed in live fish. We also discover patterns which lead to a marked increase in swimming-speed, albeit with a significant increase in energy expenditure. Notably, the use of multi-objective optimization reveals locomotion patterns that strike the perfect balance between speed and efficiency, which can be invaluable for use in robotic applications. The analyses presented may also be extended for optimal design and control of airborne vehicles. As an additional goal of the paper, we highlight the ease with which disparate codes can be coupled via the software framework used, without encumbering the user with the details of efficient parallelization.

I. INTRODUCTION

Steady, continuous swimming for long duration is rarely observed in most fish species. Instead, individuals usually employ short bursts of activity followed by a brief inactive period, where inertia allows them to glide along for a certain distance [1]. This unsteady swimming pattern, referred to as ‘burst-and-coast’ swimming, has been hypothesized to yield energetic benefits [2], [3], and to bestow a competitive edge to species employing the technique. This sort of intermittent motion is not just limited to fish-swimming, but is found in a variety of animal species [4]. In addition to reducing energy expenditure, the inactive phase of the motion has been attributed with stabilizing the sensory field, enhancing the possibility of prey-detection, and diminishing the wake-signature to avoid alerting potential prey and predators [5], [6]. Unfortunately, all of these advantages are usually accompanied by a reduction in average speed of the organism [7].

Early studies investigating the energetic benefits of intermittent swimming relied on simplified inviscid flow-dynamics [2], [3]. Certain approaches have even represented fish bodies using prolate spheres [8]. Videler & Weihs [3] analysed video of free swimming cod and saithe, and used simplified energetics models to surmise that the fish used burst-coast patterns that resulted in the most efficient locomotion. Müller et al. [9] and Videler et al. [7] used particle image velocimetry (PIV) to study flow-patterns generated during burst-coast swimming. McHenry & Lauder investigated the dependence of coasting behaviour on changes in the body-morphology [10]. Wu et al. [11] demonstrated the existence of two distinct tail-beat modes for swimming carp, which adopt burst-coast swimming for approximately 45% to 75% of their swimming time. Apart from various experimental investigations, burst-coast swimming has also been the subject of some numerical studies [12], [13], albeit using a-priori specified parameters for the burst-coast motion.

In order to discover the best burst-and-coast swimming approach, without having to resort to simplifying assumptions, we couple high-fidelity simulations of self-propelled swimmers with evolutionary-optimization algorithms [14], [15]. The use of multi-objective optimization allows us to consider two important, and potentially conflicting metrics, namely, the average speed, and the Cost of Transport (CoT). These two quantities dictate overall performance for the majority of motile organisms, and consequently, determine their evolutionary fitness to pass on advantageous characteristics to future generations. Analyzing flow-patterns generated by individuals that emerge as optimal solutions can provide invaluable insight regarding intermittent-swimming modes, which can help us improve the efficacy of underwater propulsion systems. Furthermore, we investigate the trade-off between speed and efficiency that fish experience when transitioning from the larval to the adult stage of their lives. This is essential for ascertaining the causal-effects that influence the evolution of the burst-coast pattern, as certain studies have shown that the preference for passive coasting increases with age (and hence, with the body size) [16].

The current work is distinct from previous numerical and analytical studies that have utilized optimization algorithms to examine self-propelled swimmers [17]–[22], in

that we study intermittent-locomotion modes. Moreover, we consider two disparate performance metrics simultaneously, both of which are critical for the survival of most organisms. From a computational viewpoint, the modular approach used for coupling standalone applications in this study, each of which is highly specialized for its own particular task, allows for nearly effortless, and extremely effective task-based parallelization of the optimization procedure. The parallelization is further enhanced by machine-learning based task-scheduling.

II. SIMULATION DETAILS

We use a software framework, ‘MRAG-I2D’ [23], to conduct simulations of self-propelled fish-like swimmers. The code is an open source framework for simulating two-dimensional, viscous, incompressible flows on multi-core architectures. The use of multi-resolution grids, capable of adapting automatically in both space and time, enables accurate simulation of physical systems, while keeping both computational cost and memory requirement low. The solver used in the present study is based on the remeshed vortex methods [24], and has been validated and used extensively for simulations of complex, deforming objects [18], [21], [25].

A. Governing equations

The velocity field in the simulations is governed by the incompressible Navier-Stokes equations:

$$\nabla \cdot \mathbf{u} = 0 \quad (1)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = \frac{-\nabla P}{\rho} + \nu \nabla^2 \mathbf{u} \quad (2)$$

Here, \mathbf{u} represents the velocity of the fluid, P is the pressure, ρ is the fluid-viscosity, and t is the simulation time. The operators $\nabla \cdot (\cdot)$, $\nabla (\cdot)$, and $\nabla^2 (\cdot)$ represent the divergence, the gradient, and the Laplacian, respectively. The interaction between fluid-flow and solid objects is achieved via Brinkman penalization [26], which leads to a modified form of the momentum equation:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = \frac{-\nabla P}{\rho} + \nu \nabla^2 \mathbf{u} + \lambda \chi (\mathbf{u}_s - \mathbf{u}) \quad (3)$$

Here, λ is a penalty parameter, and χ is the characteristic function describing the distribution of the solid object on a the Cartesian fluid-grid. The symbol \mathbf{u}_s in Eq. 3 denotes the pointwise velocity of the discretized solid, and accounts for translation, rotation, and deformation of the body. The vorticity (ω) form of the momentum equation is obtained by taking the curl of Eq. 3:

$$\frac{\partial \omega}{\partial t} + (\mathbf{u} \cdot \nabla) \omega = \nu \nabla^2 \omega + \lambda \nabla \times (\chi (\mathbf{u}_s - \mathbf{u})) \quad (4)$$

Further details regarding spatial discretization, fluid-solid interaction, and the time-splitting steps involved in solving Eq. 4 are described in refs. [23], [25].

B. Swimmer kinematics & Optimization parameters

The undulatory kinematics of the self-propelled swimmer is described by the following travelling wave:

$$y_{midline}(s, t) = f(t) \cdot \frac{4}{33} (s + 0.03125L) \cdot \sin \left(2\pi \left(\frac{s}{L} - \frac{t}{T} + \phi \right) \right) \quad (5)$$

Here, s is the curvilinear parameter representing the backbone of the fish from the head to the tail tip, with $s \in [0, L]$. L is the length of the fish, T denotes the time-period of the travelling wave, and ϕ represents a phase shift. Further details regarding the description of the body-geometry may be found in [25].

The function $f(t)$ controls the envelope of the time-varying amplitude of the swimmer. It is defined as a piecewise continuous function, with intervals designated by the time points t_A, t_B, t_C and t_D . At t_A , the fish initiates a deceleration by reducing its undulatory movement. The swimmer ceases all motion at t_B , and glides in an unpowered state until t_C . At t_C , the swimmer restarts its motion until reaching full amplitude at t_D . For a swimmer performing the i -th burst-and-coast cycle, the piecewise function $f(t)$ is computed as follows:

$$f(t) = \begin{cases} 1 & t_D^{(i-1)} \leq t < t_A^{(i)} \\ 1 - 3\lambda_{coast}^2 + 2\lambda_{coast}^3 & t_A^{(i)} \leq t < t_B^{(i)} \\ 0 & t_B^{(i)} \leq t < t_C^{(i)} \\ 3\lambda_{burst}^2 - 2\lambda_{burst}^3 & t_C^{(i)} \leq t < t_D^{(i)} \\ 1 & t_D^{(i)} \leq t \end{cases} \quad (6)$$

Here, $\lambda_{coast}, \lambda_{burst} \in [0, 1]$ are ramp functions increasing linearly from 0 to 1 as t goes from t_A to t_B (t_C to t_D respectively). This particular definition of the piecewise function prevents discontinuous jerks and acceleration, thereby avoiding spurious numerical oscillations in the flow-field.

The controlling parameters for the optimizations are then:

$$t_{Decel} = t_B^{(i)} - t_A^{(i)} \in [0.1, 1.5] \quad (7)$$

$$t_{Coast} = t_C^{(i)} - t_B^{(i)} \in [0, 3] \quad (8)$$

$$t_{Accel} = t_D^{(i)} - t_C^{(i)} \in [0.1, 1.5] \quad (9)$$

$$t_{Steady} = t_A^{(i+1)} - t_D^{(i)} \in [0, 2] \quad (10)$$

This choice of controlling parameters allows a simple representation of burst-and-coast swimming, and keeps the dimension of the search-space to a reasonable size. The lower bounds, $t_{Decel} \geq 0.1$ and $t_{Accel} \geq 0.1$ prevent the swimmers from experiencing excessively large accelerations.

C. Performance metrics

The objective of the self-propelled swimmers is to discover the best mode of intermittent swimming, which maximizes the average speed, while at the same time minimizing the energy consumed for travelling a unit distance (also referred to as the Cost of Transport). Computing the Cost of Transport (or CoT) requires knowledge of the pressure-

and viscosity-induced forces:

$$d\mathbf{F}_p = -P\mathbf{n} dS \quad (11)$$

$$d\mathbf{F}_v = 2\mu\mathbf{D} \cdot \mathbf{n} dS \quad (12)$$

Here, $\mathbf{D} = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$ is the strain-rate tensor, μ is the dynamic viscosity of the fluid, \mathbf{n} is the surface normal, and dS is the infinitesimal surface area. The pressure P is computed from a Poisson's equation obtained by taking the divergence of the penalized momentum equation (Eq. 3), and utilizing the incompressibility condition ($\nabla \cdot \mathbf{u} = 0$):

$$\nabla^2 P = -\rho(\nabla\mathbf{u}^T : \nabla\mathbf{u}) + \rho\lambda\nabla \cdot (\chi(\mathbf{u}_s - \mathbf{u})) \quad (13)$$

This Poisson's equation is solved via an extremely efficient tree-code algorithm, which is based on multipole expansions [27]. Moreover, pressure computations are carried out only at the surface nodes of the body, which reduces the computational cost significantly.

The thrust generated by the swimmers is determined as follows:

$$\text{Thrust} = \frac{1}{2\|\mathbf{u}\|} \iint (\mathbf{u} \cdot d\mathbf{F} + |\mathbf{u} \cdot d\mathbf{F}|) \quad (14)$$

where $d\mathbf{F} = d\mathbf{F}_p + d\mathbf{F}_v$. Using these quantities, the thrust- and deformation-power are computed as:

$$P_{\text{Thrust}} = \text{Thrust} \cdot \|\mathbf{u}\| \quad (15)$$

$$P_{\text{Def}} = - \iint \mathbf{u}_{\text{Def}} \cdot d\mathbf{F} \quad (16)$$

where \mathbf{u}_{Def} represents the deformation-velocity of the fish body. The instantaneous swimming-efficiency is based on a modified form of the Froude efficiency proposed in ref. [28]:

$$\eta = \frac{P_{\text{Thrust}}}{P_{\text{Thrust}} + \max(P_{\text{Def}}, 0)} \quad (17)$$

To compute both η and the Cost of Transport (CoT), we neglect negative values of P_{Def} , which corresponds to the fact that the rigid fish-body may not store energy internally:

$$\text{CoT}(t) = \frac{\int_{t-T_p}^t \max(P_{\text{Def}}, 0) dt}{\int_{t-T_p}^t \|\mathbf{u}\| dt} \quad (18)$$

This restriction yields a conservative estimate of potential savings in the CoT. To avoid the initial transient when the fish accelerates from rest, both the average speed and average CoT are determined over the final 2 burst-coast cycles in a simulation, as depicted in Fig. 1. These two quantities represent the 'fitness' values used for multi-objective optimization, which is described in greater detail in Section III. The Reynolds number (Re , mentioned in Fig. 1 caption) is computed as $Re = L^2/(\nu T)$, where L is the length of the swimmer's body, and T is the tail-beat period. This number indicates the relative contribution of inertial-forces to viscous-forces in the flow. Viscous-effects tend to dominate at lower Reynolds numbers, as in the case of a fish-larva ($Re = 400$), whereas inertial effects dominate for an adult-fish ($Re = 4000$).




Fig. 1: The instantaneous speed of a swimmer employing intermittent locomotion at $Re = 400$. The highlighted region depicts the averaging zone used for computing the fitness value. The averaging is done over the last two complete cycles to avoid unsteady transients towards the beginning of the simulation.

III. OPTIMIZATION ALGORITHMS

Evolutionary-optimization algorithms are designed to mimic the rules of evolution in biology, and involve the processes of mutation, reproduction, and selection. These strategies operate on a collection of individuals to iteratively improve the population with respect to a certain 'fitness' value. The fitness value represents the eligibility of a particular individual to pass on its characteristics to future generations, and in our case is determined via high-fidelity simulations of self-propelled swimmers.

We focus on multi-objective fitness functions that produce an output in \mathbb{R}^2 , which can be ordered using the concept of Pareto dominance. A vector $\mathbf{u} \in \mathbb{R}^n$ dominates a vector $\mathbf{v} \in \mathbb{R}^n$ in terms of Pareto dominance if the following two conditions are met:

$$\forall i \in \{1, \dots, n\} \quad u_i \leq v_i \quad (19)$$

$$\exists j \in \{1, \dots, n\} \quad u_j < v_j \quad (20)$$

Any vector that is not dominated by any other vector in the set is called 'non-dominated'. Each fitness value is associated with a non-domination rank, which is set to 0 for non-dominated solutions. Further non-domination ranks are computed iteratively using the fitness values, after excluding individuals with lower ranks.

The two algorithms described below use the concept of Pareto dominance to solve multi-objective optimization problems, by assigning each individual of the population a rank according to its level of non-dominance. Individuals with the same non-dominance ranks are ordered either using the S-metric (MO-CMA-ES [29]), or the crowding distance parameter (NSGA-II). The algorithm parameters, such as population size and mutation probability were not varied in the current study, owing to the relatively high computational cost of individual function-evaluations. We remark that the use of meta-models (such as simplified inviscid vortex dynamics [2], [3], or Lighthill's slender-body theory [30]) may not yield accurate results, considering the fact that optimal individuals exhibit markedly different behaviour depending on the Reynolds number (relevant discussion in Section V).

A. Multi-Objective Covariance Matrix Adaptation Evolution Strategy (MO-CMA-ES)

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [31] is one of the most powerful evolutionary algorithms for real-valued optimization. It omits the concept of reproduction, and relies solely on mutation. The MO-CMA-ES algorithm [15] implemented in the ‘Shark’ C++ library [32] was used for multi-objective optimization in the current work. An outline of the steps involved is shown in Algorithm 1, with further details described in the text below.

Algorithm 1: MO-CMA-ES

```

1 Initialize population  $P^{(0)}$  of size  $\mu$  and set generation
  counter  $g = 0$ .
2 while No stopping criterion is met do
3   for  $k = 1, \dots, \mu$  do
4     | Sample an offspring  $\tilde{x}_k \sim N(x_k^{(g)}, \sigma_k^{(g)}, C_k^{(g)})$ .
5   end
6   Choose the best  $\mu$  individuals from a mixed
  population containing parents and offspring and
  store them in  $P^{(g+1)}$ .
7   for  $k = 1, \dots, \mu$  do
8     | Update success rate  $\bar{p}_{succ,k}^{(g+1)}$  and global step size
   $\sigma_k^{(g+1)}$ .
9     | Update evolution path  $p_{k,c}^{(g+1)}$  and the covariance
  matrix  $C_k^{(g+1)}$ .
10  end
11   $g = g + 1$ 
12 end
13 Output the current population.

```

The k -th individual in the g -th generation is described by a 5-tuple:

$$a_k^{(g)} = [x_k^{(g)}, \bar{p}_{succ,k}^{(g)}, \sigma_k^{(g)}, p_{c,k}^{(g)}, C_k^{(g)}] \quad (21)$$

Here, x_k represents the vector of search-parameters, and $\bar{p}_{succ,k}$ denotes the average success rate of this individual. σ_k is the global step size, $p_{c,k}$ is the cumulative evolution path, and C_k is the covariance matrix. The average success rate $\bar{p}_{c,k}$ is a heuristic used to manage the global step size σ_k , and the covariance matrix C_k determines the most probable direction of the mutation. The covariance matrix is updated using the evolution path $p_{c,k}$, which keeps track of the mutation-direction of the previous few steps. The success rate of each individual is updated as follows:

$$\bar{p}_{succ,k}^{(g+1)} = (1 - c_p)\bar{p}_{succ,k}^{(g)} + c_p p_{succ,k}^{(g)} \quad (22)$$

where c_p is a constant weighting factor, and $p_{succ,k}^{(g)}$ is equal to one if the offspring of the k -th individual dominates its parent. The step size update then becomes:

$$\sigma_k^{(g+1)} = \sigma_k^{(g)} \cdot e^{\left(\frac{1}{d} \cdot \frac{\bar{p}_{succ,k}^{(g+1)} - p_{succ}^{target}}{1 - p_{succ}^{target}}\right)} \quad (23)$$

Here, d is a damping parameter, and p_{succ}^{target} is the target success probability. Intuitively, the step size drops if the smoothed success probability $\bar{p}_{succ,k}$ is smaller than the target success probability p_{succ}^{target} , and grows otherwise. The update of the the covariance matrix is done differently depending on whether \bar{p}_{succ} is smaller than some threshold $p_{thresh} \in [0, 1]$ ([15]). This is necessary to prevent the eigenvalues of C_k from becoming too large.

B. Non-dominated Sorting Genetic Algorithm-II (NSGA-II)

Non-dominated Sorting Genetic Algorithm II (NSGA-II) [14] improves on the original NSGA [33] by introducing elitism, the crowding distance parameter, and a fast non-dominated sorting algorithm. Elitism refers to the best individuals of a generation being carried over to the next generation. The crowding distance parameter is an estimate of the average side-length of the cuboid formed by an individual’s two closest neighbours. Individuals with a higher crowding distance are preferred because they increase diversity in the solution. The steps involved in the NSGA-II algorithm are as follows:

- 1) Initialize population $P^{(0)}$ of size μ and set generation counter $g = 0$
- 2) While no stopping criterion is met
 - Generate an offspring population $\bar{P}^{(g+1)}$ from the parent population $P^{(g)}$ using crossover (SBX) and mutation (PBM)
 - Combine offspring $\bar{P}^{(g+1)}$ and parent population $P^{(g)}$ into a mixed population and sort it according to non-domination rank and crowding distance
 - Select the best μ individuals for the next parent generation $P^{(g+1)}$
 - $g = g + 1$
- 3) Output the current population

The k -th individual in generation g is described by a vector $x_k^{(g)}$ containing the search parameters. The crossover step of the algorithm is performed using Simulated Binary Crossover (SBX), whereas Parameter-based Mutation (PBM) is used for the mutation step [34]. All the individuals are represented by real-valued, n -dimensional vectors, where n is the dimension of the search-space (i.e., the number of parameters being optimized). In the current study, the probability of crossover was set to 0.9, with a crossover distribution index of 5, whereas the probability of mutation was set to 0.5 with a mutation distribution index of 10.

IV. HIGH PERFORMANCE OPTIMIZATION FRAMEWORK

The parallel implementation of the optimization algorithms is based on the TORC task-parallel library [35]. TORC provides a programming and runtime environment similar to OpenMP tasks, but allows parallel programs to run on both shared and distributed memory systems. MPI applications run on the compute nodes with one or more workers, and can submit tasks for asynchronous execution from any nesting level of parallelism. The library exports

a C/C++ and Fortran interface, provides transparent task and data management, and supports load balancing through intra- and inter-node work stealing, enhanced by a set of task distribution policies that can be applied by the programmer.


TORC has already been used for the parallelization of algorithms for single objective stochastic optimization and uncertainty quantification algorithms, as part of the $\Pi 4U$ [36] framework. It has been shown to achieve excellent parallel efficiency on up to 512 compute nodes of Piz Daint, even in cases where the simulation time exhibits significant variance depending on the input parameters. We note that only the function-evaluations are parallelized in the current work, whereas the optimization algorithm, which is executed only intermittently upon completion of a generation, is run in serial on a master thread. Function-evaluations correspond to multi-core simulations performed by MRAG-I2D, using input parameters specified by the optimization algorithm. We utilize a single worker per compute node, launch MRAG-I2D with the fork-exec system calls, and perform data exchange through the local filesystem.

The large variance of the MRAG-I2D simulation-time introduces load imbalance that can lead to poor hardware utilization. This can be avoided by setting the population size of the optimization algorithms to be larger than the number of available workers, and exploiting the task-stealing mechanism of the library. To further assist this approach, at every optimization step, function-evaluation tasks with larger expected runtime are scheduled first. In particular, we sort the tasks in descending order, distribute cyclically the first W of them to the W available workers, and insert the remaining to a single queue which is checked by idle workers for pending tasks.


To estimate the expected runtime for each function evaluation of a new generation, we exploit the information available from all previous generations. More specifically, we design a fully connected artificial neural network, with one hidden layer consisting of 32 neurons and a non-saturating activation function. The training dataset consists of the simulation parameters (input) and the corresponding estimated runtimes (output). We optimize the mean squared error of the network using the L-BFGS optimization algorithm. The open source OpenANN library¹ was used for the implementation of the neural network. We note that managed scheduling in this manner can yield up to a 20% increase in parallel efficiency for the optimization campaigns.

V. RESULTS

The optimization studies were conducted at two different Reynolds numbers, $Re = 400$ and $Re = 4000$, in order to characterize differences in the swimming behaviour of fish larvae and adults. A single optimization campaign required approximately 20 generations consisting of 32 individuals



(a)



(b)

Fig. 2: Pareto front obtained from the optimization of burst-and-coast swimmers at (a) $Re = 400$, and (b) $Re = 4000$. The front obtained from the NSGA-II algorithm is shown in blue, whereas the front obtained using MO-CMA-ES is shown in red. The dominated individuals (non-optimal) are depicted as open circles.

each. Multiple distinct simulations were run in parallel, with each simulation representing a single individual of the population, and utilizing 8 threads (spanning an entire compute node) to average a wall-clock time of approximately 3 hours. The lower bounds for the 4 optimization-parameters introduced in Section II-B were specified to be $[0.1, 0, 0.1, 0]$, whereas the upper bounds were set to $[1.5, 3.0, 1.5, 1.5]$. At the end of the optimization runs, the individuals that comprise the best possible solutions from the entire population constitute the ‘Pareto front’, as shown in Fig. 2. The figure also depicts a variety of non-optimal individuals that were obtained as part of the optimization process. The entire population shown in Figs. 2a and 2b is comprised of individuals obtained from both the MO-CMA-ES and the NSGA-II algorithms. We observe that for the

¹<https://github.com/OpenANN/OpenANN>

selected algorithm settings, MO-CMA-ES is able to generate a much more diverse Pareto front than NSGA-II. However, the latter may discover individuals with better fitness values in some instances (e.g., the fastest individual in Fig. 2b). We remark that all of the individuals that lie on the front represent Pareto-optimal solutions with different physical characteristics, but they may not necessarily represent the “true” underlying front. A select few of the non-dominated individuals from the combined population of both MO-CMA-ES and NSGA-II were chosen for further analysis, namely, the most efficient, a generalist, and the fastest swimmer. Of these three, the efficient and the fast swimmers were chosen as the individuals on either extreme of the Pareto front, whereas the generalist swimmer was selected as the median of the Pareto-optimal individuals.

A. Vorticity field and burst-coast characteristics

The flow-fields generated by an adult steady swimmer, and by three different adult individuals obtained as outcomes of the optimization procedure, are shown in Fig. 3. In a given time, the swimmers cover varying distances owing to differences in their average speed, which is a consequence of adopting different swimming patterns. There are notable differences in the wake-vortices, influenced by the particular intermittent modes selected by the different swimmers. From Fig. 3a, we observe that the steady swimmer generates regularly spaced wake vortices, owing to its sinusoidally varying body undulations. The efficient adult swimmer (Figs. 3b-3e, Movie 1) undergoes a long coasting period (t_{Coast}), which results in minimal vortex-shedding into the surrounding flow-field. Upon resuming body-undulations, the swimmer sheds three distinct vortices (Fig. 3e), which resemble the trio of vortices generated during the initial asymmetric start-up. The ‘generalist swimmer’, i.e., the optimal individual that embodies a useful compromise between the two conflicting objectives (i.e., high speed and high energy efficiency), spends a measurable amount of time swimming steadily, and experiences relatively gentle accelerations. The resulting wake vortices resemble a combination of those generated by the steady (Fig. 3a) and the most efficient swimmers. The fast adult swimmer, on the other hand, utilizes extremely short speed-up and slow-down times, and correspondingly large accelerations. The relevant tail-beat motion allows the fast swimmer to rapidly accelerate packets of fluid backward, which in turn generates large forward thrust. Moreover, the acceleration phase is followed by a short coasting phase, which allows the adult to move out of the unsteady transient it generates, before initiating another acceleration cycle. The resulting strong vortices that emerge in the wake are visible in Figs. 3c-3e.

Analyzing the burst-coast cycle-time quantitatively in Fig. 4, we realize that the most efficient swimmers spend a negligible amount of time swimming with a steady gait, irrespective of the developmental stage they may be in (i.e., larva or adult). This correlates well with experimental




Fig. 4: Breakdown of the burst-coast cycle time for the intermittent (a) larval swimmers ($Re = 400$), and (b) adult swimmers ($Re = 4000$).

observations, where fish executing burst-and-coast patterns are rarely observed to employ steady motion. We notice striking similarities between the burst-coast time-profiles for both the adult and larval efficient swimmers; there is an initial deceleration, followed by more than 50% of the time spent coasting, which is followed by an acceleration phase with almost no steady swimming in between. The generalist adult swimmer spends most of its time executing active motion, and coasts only for a short duration (Fig. 4). The burst-coast profile of the generalist larva is similar in nature, with the larva also adopting a short coasting stage.

The most notable differences appear in the burst-coast time-profiles of the fast adult and larval swimmers; while both of these adopt very short speed-up and slow-down times (t_{Accel} and t_{Decel}), giving rise to correspondingly large accelerations, the larva prefers to spend the majority of its time executing steady swimming, whereas the adult swimmer exhibits a strong preference for coasting. This is an important result, since the only difference in the two cases is the Reynolds number. The reluctance of the larva to undergo coasting may reflect the fact that it experiences large deceleration during coasting, due to high viscous drag, which is not the case for the adult swimmer. It is crucial that the optimizer is able to discern this key physical difference, which has a purely hydrodynamic origin. Furthermore, the



(a)

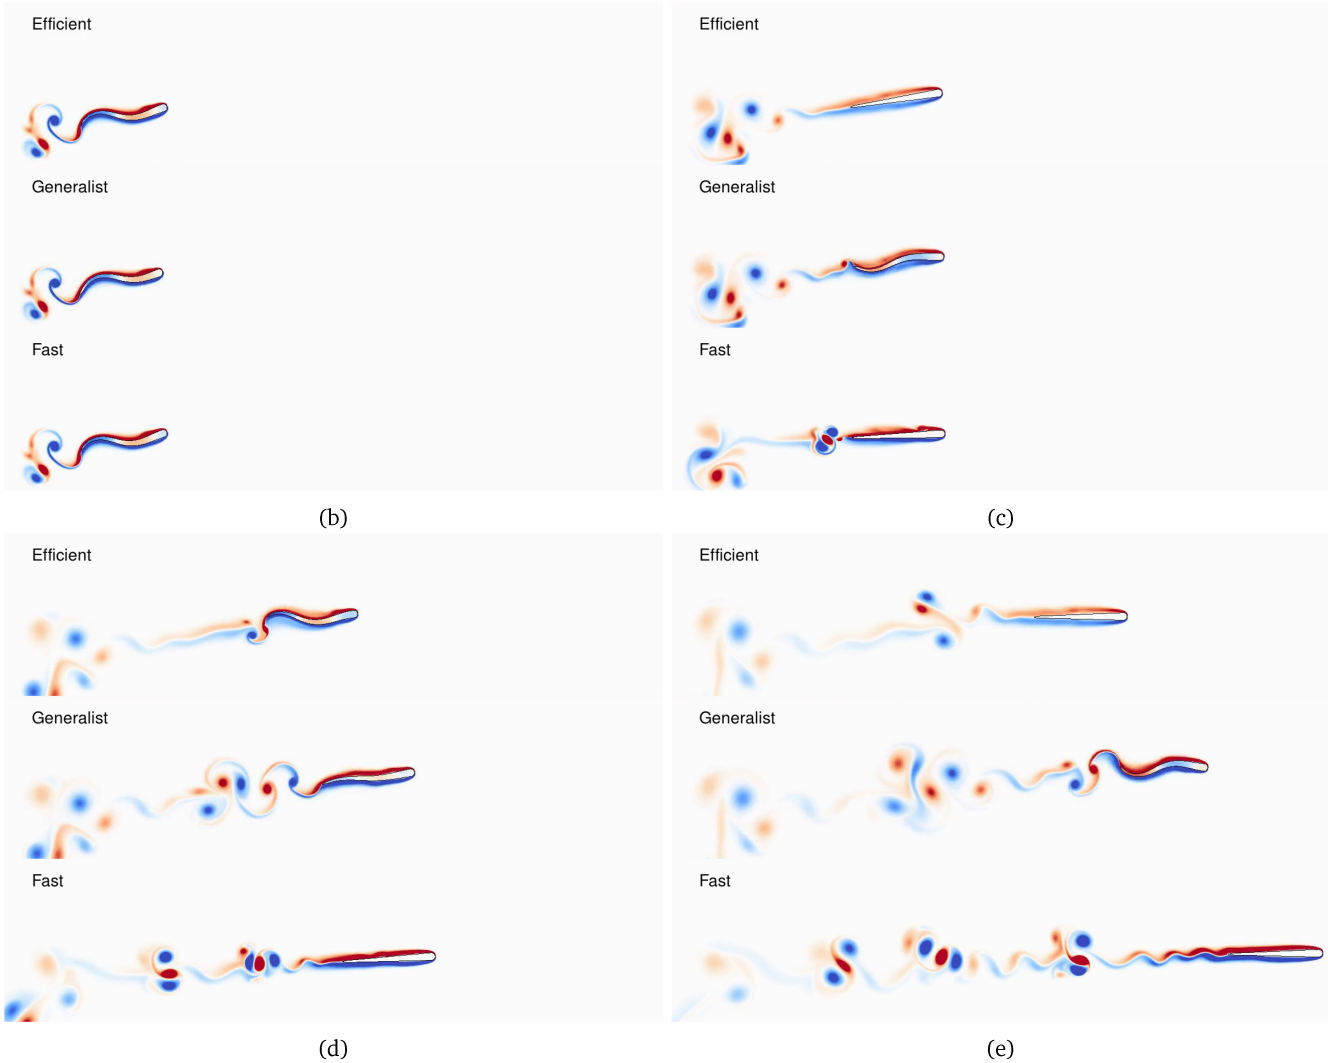


Fig. 3: The vorticity field generated by (a) a steady adult swimmer; (b) the efficient, the generalist, and the fast swimmers at $t = 2.0$ (right before the initiation of burst-coast cycles); (c) at $t = 4.5$; (d) at $t = 7.0$; and (e) at $t = 9.5$. Movie [1](#) shows an animation of the flow-fields generated by the adult swimmers. Positive vorticity, which indicates regions of fluid ‘swirling’ in an anti-clockwise manner are coloured in red, whereas negative vorticity (depicting clockwise swirl) is shown in blue. Snapshots of the flow-fields generated by the larval swimmers are not shown here, but the relevant animation may be found in Movie [2](#).

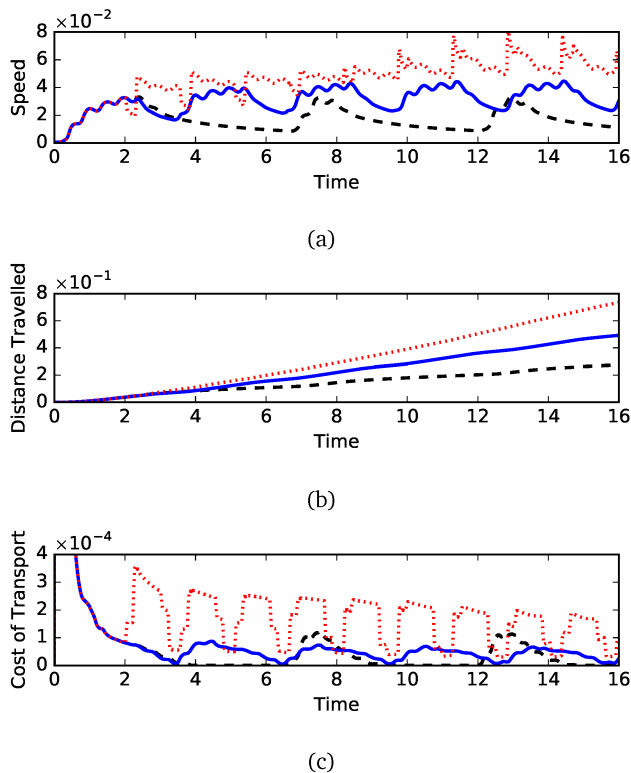


Fig. 5: (a) Speed, (b) distance travelled, and (c) cost of transport for three different larval swimmers ($Re = 400$). The dashed black line corresponds to the efficient swimmer, the solid blue line to the generalist swimmer, and the dotted red line to the fast swimmer.

propensity of both the fast larva and the adult to accelerate a significant mass of water in a short time is reminiscent of the strategy used by fish executing escape manoeuvres in life-threatening situations [18], [37], [38].

B. Performance metrics

The performance metrics of interest, namely, the speed, the distance traversed, and the Cost of Transport (CoT, Eq. 18) are shown in Fig. 5 for the larval swimmer, and in Fig. 6 for the adult swimmer. In both situations, we clearly observe the increasing trend in instantaneous speed and net distance traversed, as we move from the efficient, to the generalist, to the fast swimmer. For the sake of clarity, the metrics for the fast adult swimmer have been plotted using a different y-range, shown on the right hand side of the graphs (Fig. 6). Importantly, the peak CoT values (Figs. 5c and 6c) show a marked increase in the case of the fast swimmers (approximately $2\times$ for the larva, and $8\times$ for the adult), but are comparable for the efficient and the generalist swimmers. The peak power output is usually constrained by the musculature of an organism, in addition to internal metabolic limitations, which may explain why the fast intermittent kinematics observed in Movie 1 is generally not encountered in nature [20].

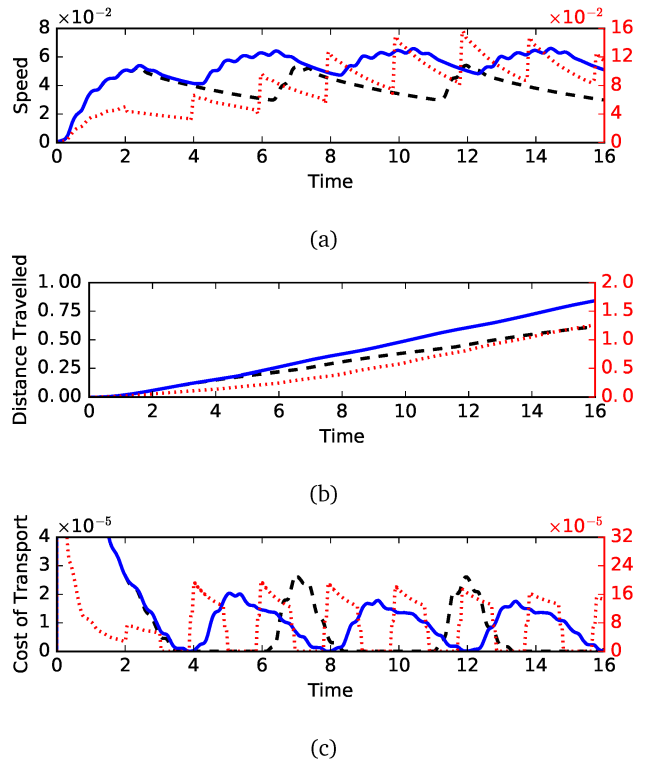


Fig. 6: (a) Speed, (b) distance travelled, and (c) cost of transport for three different adult swimmers ($Re = 4000$). The dashed black line corresponds to the efficient swimmer, the solid blue line to the generalist swimmer, and the dotted red line to the fast swimmer. The secondary y-axes shown on the right hand side of each figure correspond to the fast adult swimmer.

TABLE I: Fitness data for larval swimmers at $Re = 400$. The values have been normalized with respect to the data for the Steady swimmer.

Swimmer	CoT	Avg. Speed
Steady	1.00	1.00
Efficient	0.62	0.53
Generalist	0.87	0.84
Fast	2.70	1.47

TABLE II: Fitness data for adult swimmers at $Re = 4000$. The values have been normalized with respect to the data for the Steady swimmer.

Swimmer	CoT	Avg. Speed
Steady	1.00	1.00
Efficient	0.48	0.63
Generalist	0.77	0.88
Fast	3.57	1.95

The fitness values, computed by averaging the appropriate quantities over the last two burst-coast cycles, are listed in Tables I and II for the larval and adult swimmers, respectively. Towards the end of a swimming bout (marked by the termination time $t = 16$ for all cases), the fast larva

experiences close to a 47% increase in average speed compared to a steadily-swimming larva (Table I). However, this increase in speed is accompanied by a 170% increase in the amount of energy spent for covering a unit distance (i.e., the CoT). On the other hand, the burst-coast mode allows the most efficient larva to reduce energy-consumption by 38%, but with a reduction of 47% in average speed. Similarly, the generalist larva reduces energy consumption by 13% and average speed by 16%.

The benefits of burst-coast swimming are more pronounced in the case of the adult swimmer, as can be surmised by examining the data in Table II. The most efficient adult reduces energy consumption by 52%, but with only a 37% drop in average speed. The generalist adult swimmer shows excellent performance gains, and reduces energy expenditure by almost a quarter, with merely a 12% reduction in average speed. The fast adult swimmer sees a 2× speed-up compared to a steady swimmer, but with an immense increase (257%) in CoT. Moreover, the adult is capable of attaining much higher speeds than the larva, at the expense of much lower CoT, which supports the hypothesis that rapid increase of body-length is a high priority for fish larvae, so as to escape the detrimental effects of high viscosity [39]. These results suggest that it is feasible to discover a diverse array of optimal locomotion techniques by combining high-fidelity simulations with multi-objective optimization techniques, especially when conflicting objectives may need to be considered simultaneously. Future investigations will build upon the current work by examining the optimal swimmers' motion from a hydrodynamic perspective, with the aim of better understanding the physics driving the evolution of intermittent locomotion in fish-like swimmers.

VI. CONCLUSION

We have coupled task-parallel multi-objective optimization algorithms with high-fidelity simulations of self-propelled swimmers, to discover efficient intermittent locomotion modes in self-propelled swimmers. This approach has allowed us to discover a diverse range of motion patterns, each of which are optimal in their own right, but give rise to very different locomotion characteristics. The most efficient swimmers discovered by the optimization algorithms display prolonged coasting behaviour, which allows them to conserve energy. The fastest swimmers were able to generate large forward thrust by rapidly accelerating 'packets' of fluid opposite to their direction of motion. Intriguingly, the fast larva and the fast adult adopt markedly differently burst-coast behaviour, which may be ascribed to the increased viscous drag experienced by the larva during coasting. This indicates that the optimization algorithms are capable of correctly accounting for expected differences in the physics of the problem. The primary advantage of using such computational tools lies in the fact that it can lead to the discovery of patterns which may not generally occur in nature, but can be invaluable for use in robotic applications; optimal swimming gaits can be selected specifically to suit

particular mission-requirements, e.g., to maximize energy conservation, minimize travel time, or utilize an optimal combination of both. The resulting increase in range, endurance, and average speed can greatly enhance the mission capability of robotic swimmers. The tools and techniques described here are not only relevant for studying artificial swimmers, but may also be used for design- and control-optimization of a variety of aerial vehicles.

ACKNOWLEDGMENTS

This work was supported by the European Research Council Advanced Investigator Award, and the Swiss National Science Foundation Sinergia Award (CRSII3 147675). Computational resources were provided by the Swiss National Supercomputing Center (CSCS) under project ID 's658'.

REFERENCES

- [1] L. A. Fuiman and P. W. Webb, "Ontogeny of routine swimming activity and performance in zebra danios (teleostei: Cyprinidae)," *Anim. Behav.*, vol. 36, no. 1, pp. 250 – 261, 1988.
- [2] D. Weihs, "Energetic advantages of burst swimming of fish," *J. Theor. Biol.*, vol. 48, no. 1, pp. 215–229, 1974.
- [3] J. J. Videler and D. Weihs, "Energetic advantages of burst-and-coast swimming of fish at high speeds," *J. Exp. Biol.*, vol. 97, no. 1, pp. 169–178, 1982.
- [4] P. Paoletti and L. Mahadevan, "Intermittent locomotion as an optimal control strategy," *Proc. Royal Soc. London A: Math. Phys. Eng. Sci.*, vol. 470, no. 2164, 2014.
- [5] D. L. Kramer and R. L. McLaughlin, "The behavioral ecology of intermittent locomotion," *Amer. Zool.*, vol. 41, no. 2, pp. 137–153, 2001.
- [6] S. Wieskotten, G. Dehnhardt, B. Mauck, L. Miersch, and W. Hanke, "The impact of glide phases on the trackability of hydrodynamic trails in harbour seals (phoca vitulina)," *J. Exp. Biol.*, vol. 213, no. 21, pp. 3734–3740, 2010.
- [7] J. J. Videler, E. J. Stamhuis, U. K. MÅijller, and L. A. van Duren, "The scaling and structure of aquatic animal wakes," *Integr. Comp. Biol.*, vol. 42, no. 5, pp. 988–996, 2002.
- [8] R. W. Blake, "Functional design and burst-and-coast swimming in fishes," *Canadian Journal of Zoology*, vol. 61, no. 11, pp. 2491–2494, 1983.
- [9] U. K. Muller, E. J. Stamhuis, and J. J. Videler, "Hydrodynamics of unsteady fish swimming and the effects of body size: comparing the flow fields of fish larvae and adults," *J. Exp. Biol.*, vol. 203, no. 2, pp. 193–206, 2000.
- [10] M. J. McHenry and G. V. Lauder, "The mechanical scaling of coasting in zebrafish (danio rerio)," *J. Exp. Biol.*, vol. 208, no. 12, pp. 2289–2301, 2005.
- [11] G. Wu, Y. Yang, and L. Zeng, "Kinematics, hydrodynamics and energetic advantages of burst-and-coast swimming of koi carps (cyprinus carpio koi)," *J. Exp. Biol.*, vol. 210, no. 12, pp. 2181–2191, 2007.
- [12] M.-H. Chung, "On burst-and-coast swimming performance in fish-like locomotion," *Bioinspir. Biomim.*, vol. 4, no. 3, 2009.
- [13] M. Bergmann and A. Iollo, "Modeling and simulation of fish-like swimming," *J. Comput. Phys.*, vol. 230, pp. 329 – 348, 2011.
- [14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multi-objective genetic algorithm: Nsga-ii," 2000.
- [15] C. Igel, T. Suttorp, and N. Hansen, *Steady-State Selection and Efficient Covariance Matrix Update in the Multi-objective CMA-ES*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 171–185.
- [16] U. K. Müller, J. L. van Leeuwen, S. van Duin, and H. Liu, "An unmomentous start to life: Can hydrodynamics explain why fish larvae change swimming style?" *J. Biomech. Sci. Eng.*, vol. 4, no. 1, pp. 37–53, 2009.
- [17] S. Kern and P. Koumoutsakos, "Simulations of optimized anguilliform swimming," *J. Exp. Biol.*, vol. 209, pp. 4841–4857, 2006.
- [18] M. Gazzola, W. M. van Rees, and P. Koumoutsakos, "C-start: optimal start of larval fish," *Journal of Fluid Mechanics*, vol. 698, pp. 5–18, May 2012.

- [19] G. Tokić and D. K. P. Yue, "Optimal shape and motion of undulatory swimming organisms," *Proc. Roy. Soc. London B: Biol. Sci.*, 2012.
- [20] C. Eloy, "On the best design for undulatory swimming," *J. Fluid Mech.*, vol. 717, pp. 48–89, 002 2013.
- [21] W. M. van Rees, M. Gazzola, and P. Koumoutsakos, "Optimal shapes for anguilliform swimmers at intermediate reynolds numbers," *J. Fluid Mech.*, vol. 722, 5 2013.
- [22] —, "Optimal morphokinematics for undulatory swimmers at intermediate reynolds numbers," *J. Fluid Mech.*, vol. 775, pp. 178–188, 007 2015.
- [23] D. Rossinelli, B. Hejazialhosseini, W. M. van Rees, M. Gazzola, M. Bergdorf, and P. Koumoutsakos, "MRAG-I2D: Multi-resolution adapted grids for remeshed vortex methods on multicore architectures," *J. Comput. Phys.*, vol. 288, pp. 1–18, 2015.
- [24] P. Koumoutsakos and A. Leonard, "High-resolution simulations of the flow around an impulsively started cylinder using vortex methods," *J. Fluid Mech.*, vol. 296, pp. 1–38, 1995.
- [25] M. Gazzola, P. Chatelain, W. M. van Rees, and P. Koumoutsakos, "Simulations of single and multiple swimmers with non-divergence free deforming geometries," *J. Comput. Phys.*, vol. 230, pp. 7093–7114, 2011.
- [26] P. Angot, C. H. Bruneau, and P. Fabrie, "A penalization method to take into account obstacles in incompressible viscous flows," *Numer. Math.*, vol. 81, pp. 497–520, 1999.
- [27] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.*, vol. 73, pp. 325–348, 1987.
- [28] E. D. Tytell and G. V. Lauder, "The hydrodynamics of eel swimming," *J. Exp. Biol.*, vol. 207, pp. 1825–1841, 2004.
- [29] C. Igel, N. Hansen, and S. Roth, "Covariance matrix adaptation for multi-objective optimization," *Evol. Comput.*, vol. 15, no. 1, pp. 1–28, 2007.
- [30] M. J. Lighthill, "Note on the swimming of slender fish," *J. Fluid Mech.*, vol. 9, pp. 305–317, 1960.
- [31] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evol. Comput.*, vol. 9, no. 2, pp. 159–195, Jun. 2001.
- [32] C. Igel, V. Heidrich-Meisner, and T. Glasmachers, "Shark," *Journal of Machine Learning Research*, vol. 9, pp. 993–996, 2008.
- [33] N. Srinivas and K. Deb, "Multiobjective optimization using nondominated sorting in genetic algorithms," *Evol. Comput.*, vol. 2, no. 3, pp. 221–248, Sep. 1994.
- [34] K. Deb and S. Agrawal, *A Niche-Penalty Approach for Constraint Handling in Genetic Algorithms*. Vienna: Springer Vienna, 1999, pp. 235–243.
- [35] P. E. Hadjidoukas, E. Lappas, and V. V. Dimakopoulos, "A runtime library for platform-independent task parallelism," in *2012 20th Euro-micro International Conference on Parallel, Distributed and Network-based Processing*. IEEE, 2012, pp. 229–236.
- [36] P. E. Hadjidoukas, P. Angelikopoulos, L. Kulakova, C. Papadimitriou, and P. Koumoutsakos, "Exploiting task-based parallelism in bayesian uncertainty quantification," in *European Conference on Parallel Processing*. Springer, 2015, pp. 532–544.
- [37] P. Domenici and R. Blake, "The kinematics and performance of fish fast-start swimming," *Journal of Experimental Biology*, vol. 200, no. 8, pp. 1165–1178, 1997.
- [38] U. K. Müller, J. G. M. van den Boogaart, and J. L. van Leeuwen, "Flow patterns of larval fish: undulatory swimming in the intermediate flow regime," *Journal of Experimental Biology*, vol. 211, no. 2, pp. 196–205, 2008.
- [39] U. K. Müller and J. J. Videler, "Inertia as a 'safe harbour': do fish larvae increase length growth to escape viscous drag?" *Rev. Fish Biol. Fish.*, vol. 6, no. 3, pp. 353–360, 1996.