Energy-Optimal Trajectory Planning for Autonomous Underwater Vehicles Using Pontryagin's Maximum Principle

Raymond Young, F. Hunter Akins, Sophia T. Merrifield

Abstract—**Pontryagin's Maximum Principle, a classical tool from optimal control theory, is used to compute energy-optimal trajectories for an autonomous underwater vehicle (AUV) navigating in deterministic, complex flows. This approach reformulates the optimal control problem as a two-point boundary value problem (TPBVP) for ordinary differential equations (ODEs) and offers computational benefits compared to alternative solutions which require the solution of partial differential equations (PDEs). This study introduces a numerical framework using pure shooting to solve the associated TPBVP and demonstrates its efficacy for trajectory optimization through analytic flows and current fields from high-resolution ocean simulations. The method is tested on two case studies: 1) a multi-day transit through the Southern California Bight, where ocean currents are dominated by mesoscale eddies and tidal variability, and energy-optimal planning predicts over 20% of savings relative to a straight-line transit, and 2) a multi-hour transit off the coast of Ram Head, St. John, USVI, where the flow is dominated by energetic tidal vortices and energy consumption varies by up to 50% within a 6 hr launch window. In both cases, the tradeoff between time and energy is analyzed for a wide range of AUVs that are parameterized by their nominal speed through water, between 0.25–2.0 m/s.**

Index Terms—**Autonomous underwater vehicles, motion planning, dynamics, optimal control, numerical methods**

I. Introduction

As Autonomous Underwater Vehicles (AUVs) have become more advanced, solving optimal control problems (OCPs) has become increasingly important in marine robotics. Early vehicle development focused on low-level control of fins, thrusters, and buoyancy systems to ensure reliable operation. However, as the technology has matured and become widely adopted, research has shifted towards higher-level autonomy to improve efficiency, safety, and scalability in large vehicle fleets [1]–[4]. This paper focuses on environmentally aware trajectory planning, a subset of AUV autonomy that intersects with optimal control. The overarching goal is to integrate environmental models and forecasts to enhance operational efficiency, ultimately enabling more effective and sustainable AUV missions. Specifically, we are interested in optimizing energy-efficient transits—a critical challenge in AUV operations due to onboard power limitations—while considering the effects of ocean currents, which significantly impact speed, energy consumption, and navigability.

An environmentally aware planning problem is defined by a starting location, an ending location, a background flow field, and an objective function which assigns a cost or reward to each feasible trajectory. Its solution is a control sequence that guides the AUV from start to finish in a manner that minimizes the cost or maximizes the reward. While we focus on the marine domain, planning in a flow field is also applicable to aerial vehicles navigating through strong winds [5]. This is a classic problem in optimal control, particularly in a continuous domain, which allows us to adapt established methods to the problem at hand.

The solutions to OCPs fall into two different frameworks. The first involves solving partial differential equations (PDEs), such as the Hamilton-Jacobi-Bellman (HJB) equation [6], which provides a necessary and sufficient condition for globally optimal solutions. The second framework employs ordinary differential equations (ODEs), often derived from the calculus of variations, to establish first-order necessary conditions for optimal trajectories in the form of a two-point boundary value problem (TPBVP). Specifically, Pontryagin's Maximum Principle (PMP) [7] provides a general framework to derive optimality conditions as a TPBVP for

R. Young, F. Hunter Akins, and S. T. Merrifield are with Scripps Institution of Oceanography, UC San Diego

a Hamiltonian system of ODEs. A trajectory that satisfies PMP is guaranteed to be locally optimal (in the sense of a functional perturbation) and can thus be considered a candidate for global optimality in the OCP.

The mathematical relationship between these two approaches is well-documented [8]–[10], and can be summarized as follows: the trajectories derived using PMP correspond to the characteristics of the HJB equation for the same OCP. However, practical considerations for their numerical implementation, respectively, highlight a trade-off between accuracy, reliability, and computational complexity. While PDE-based methods guarantee global optimality, they are computationally intensive and suffer from the curse of dimensionality. In contrast, PMP-based approaches, while sensitive to initial conditions and more prone to local convergence challenges, are often more accurate and computationally efficient [10], [11].

PDE-based methods have been extensively studied for AUV planning. One common approach, established by Lolla et al. [12], applies the theory of level sets. First derived for time-optimal paths in a 2-D environment, this methodology solves for the *reachability front*, which represents the set of all states a vehicle can reach within a given time. Once the terminal state is reached, the trajectory is computed by tracing the evolution of the level set backward to the initial state. These methods have been applied to 2-D problems [13]–[16], and more recently have been extended to 3-D environments [17], energy-optimal planning [18], [19], and Pareto-optimal trajectories that balance time and energy [20]. Alternatively, Brandman and Olson [21] provide a proof for time-optimal trajectories based on the dynamic programming principle. Their method is equivalent to the level-set method when the AUV travels faster than or equal to the background flow. The equations governing the evolution of level sets are Hamilton-Jacobi PDEs, which require specialized numerical methods [22], [23].

ODE-based methods have received less attention, and existing studies are limited to 2-D, static environments. Biferale et al. [24] employ a system of equations based on Hamiltonian dynamics as a benchmark to compare reinforcement learning approaches within a single chaotic flow field. Davis et al. [25] use the calculus of variations and Fermat's principle of least time to derive optimal trajectories, drawing an analogy to the propagation of light through varying media. For flow-field navigation in other contexts, Liebchen and Lowen [26] also apply Fermat's principle to investigate optimal trajectories for microswimmers, with applications in medical fields. While they explore various time- and energy-based cost functions, their work is limited to analytically-specified static flows.

Our paper demonstrates the efficacy of ODE-based planners by applying PMP to solve for minimum energy AUV trajectories. The method is derived for 3-D, time-varying flows and demonstrated in realistic 2-D (fixed-depth), dynamic ocean models. Of the existing ODE-based methods, the Hamiltonian formalism used in [24] is most similar to our derivation presented in Section III. However, PMP allows for a direct extension to time-varying environments which has not been analyzed in previous work. Additionally, analysis in Biferale et al. emphasizes the instability of trajectories when used as an open-loop control policy [24]. Our work focuses on a reliable numerical solution to these equations so that they can be used in a pre-planning context. PMP has also been applied to an analogous problem in the aerial domain; Dobrokhodov et al. [27] studied energy-optimal path planning for long-duration aerial drones navigating dynamic wind fields. Pereria and Silvio Gama [28] applied PMP for AUVs in the context of multi-process control. This study considered the current field itself to be controllable, and focused on an analytical statement of the OCP without numerical solutions.

Although this work focuses on the trajectory optimization in the continuous domain, the problem can also be solved in a discretized domain [29], [30] which lends itself to solutions with graph search methods such as A* or Dijkstra's algorithm [31], [32]. Many studies have applied these discrete schemes to 2-D problems involving time- and energy-optimal trajectories [29], [33], [34]. More recent work by Kularatne et al. [35], [36] extends the graph-based approach to explicitly account for time-dependent flow fields. While discrete methods can be sufficient for static environments, they, like the PDE-based methods, suffer from the curse of dimensionality. Both types of methods are solved over the entire domain, and therefore computational complexity scales with the domain size, and is impacted by fine discretization which is required to resolve complex flows. In contrast, PMP provides necessary conditions for optimality that are only required to be satisfied locally, along an individual trajectory. These conditions do not grow with the size of the domain or complexity of the current field. We propose a numerical procedure to solve the TPBVP using single shooting which transcribes the optimization problem into a root-finding problem with a minimum of d optimization parameters, where d is the dimension of the environment. Therefore, the size of the problem only depends on dimension.

The rest of the paper is organized as follows: Section II defines the dynamics, control system, and energy consumption models of a simplified AUV for which we will compute optimal trajectories. Section III presents our methodology by first applying PMP to the stated OCP, and then detailing the numerical procedure used to solve the resulting TPBVP. Section IV demonstrates solutions in increasingly complex canonical flow fields, and provides a computational comparison to the level-set for time-optimal trajectories in a dynamic, 3-D flow. Section V presents two case studies that demonstrate the capability of our methodology to compute energy-optimal trajectories through realistic, dynamic geophysical flows from high-resolution model outputs. Finally, Section VI provides a summary of the results, a discussion of the method, and directions for future work.

II. Modeling

In this section, we describe the dynamics and energy-consumption model for an AUV. We begin by presenting the vehicle's dynamics, which govern its motion in the presence of ocean currents. Next, we introduce an energy consumption model that considers both time-varying (speed-dependent) and constant components. Finally, we formally state the OCP, specifying the objective function and constraints used to determine the optimal control strategy.

A. Dynamics

The vehicle is modeled as an active tracer. Its state, $\mathbf{x}(t) : \mathbb{R} \mapsto \mathbb{R}^d$, represents the vehicle's geographical position as a function of time. The control, $\mathbf{v}(t) : \mathbb{R} \mapsto \mathcal{V}$, denotes its nominal velocity through water, and the ocean currents $\mathbf{u}(t, \mathbf{x}) : (\mathbb{R} \times \mathbb{R}^d) \mapsto \mathbb{R}^d$ are deterministic. Both two- and three-dimensional environments are considered ($d \in [2, 3]$).

The vehicle's dynamics are given by:

$$\dot{\mathbf{x}}(t, \mathbf{x}, \mathbf{v}) = \mathbf{v}(t) + \mathbf{u}(t, \mathbf{x}), \tag{1}$$

representing its velocity in a global reference frame. The control set is defined as:

$$\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^d \,|\, v_{\text{min}} \leq \|\mathbf{v}\| \leq v_{\text{max}}\}, \tag{2}$$

where v_{min} is the minimum speed required to maintain control surfaces, and v_{max} is the maximum speed determined by the vehicle's capabilities. This model is commonly used in marine planning, where the vehicle size is much smaller than the dominant environmental scales [20]. It is assumed that the low-level control over actuators, such as the propeller or fins, is handled by an onboard controller that is capable of following a prescribed course.

For clarity, we define auxiliary variables for the vehicle's orientation: θ (heading) and ϕ (pitch). The unit vector $\hat{\mathbf{v}}$ aligned with the vehicle's orientation[1] is given by:

$$\hat{\mathbf{v}} \in \mathbb{S}^{d-1} := \begin{cases} \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}, & d = 2 \\ \begin{bmatrix} \cos\theta\cos\phi \\ \sin\theta\cos\phi \\ \sin\phi \end{bmatrix}, & d = 3 \end{cases} \tag{3}$$

where $\mathbb{S}^{d-1}$ is the $(d-1)$-dimensional unit-sphere embedded in $\mathbb{R}^d$.

B. Energy Consumption

The energy consumption model consists of two terms: one accounting for cubic drag, scaled by the *dynamic power coefficient* α, and another representing a constant *hotel load* β:

$$\dot{E}(t) = \alpha\|\mathbf{v}(t)\|^3 + \beta. \tag{4}$$

This model was derived and validated for a Remote Environmental Monitoring UnitS (REMUS, *Huntington Ingalls Industries, Newport News, VA*) AUV [37], [38].

The dynamic power coefficient

$$\alpha = \frac{C_p \rho d^2}{J^3}, \tag{5}$$

[1] The vehicle's orientation is not necessarily aligned with the direction of motion, $\dot{\mathbf{x}}$, which also includes advection by local currents.

depends on known parameters for propeller diameter d and water density ρ, and calibrated values for the advance ratio J and power coefficient C_p. The advance ratio J relates propeller diameter and rotational speed to forward speed through water and is estimated during routine operations from propeller counts and distance traveled. The power coefficient C_p is determined by assuming a level flight equilibrium and balancing thrust produced by the propeller and hull drag. A cubic power law is often assumed for propeller-driven AUVs; in addition to REMUS studies, it was used in analyses of the Tethys Long-Range AUV [39]. However, trajectory planning for other autonomous marine vehicles may require more complex modeling. For example, the speed and power consumption for underwater gliders depend on angle of attack and buoyancy [40], [41], while environmentally propelled vehicles depend on external forcing such as wind or waves [42]–[44].

The hotel load β accounts for power consumption unrelated to propulsion, including command-and-control systems and sensing payloads. We assume β is constant and determined by the vehicle configuration. This assumption is reasonable for long transits when onboard systems are predictable and draw approximately steady average power. However, different operating modes, such as sense-and-react adaptive sampling, can produce strongly time-varying sensor duty cycles and communication loads, in which case a time-dependent hotel load may be required. Globally optimal path-planning with respect to cubic drag was studied by Doshi et al. [45], though the explicit inclusion of hotel load was not addressed.

C. Problem Statement

Now we formally define the trajectory optimization problem. Consider a vehicle tasked with transiting between two specified locations, $\mathbf{x}_0^*$ and $\mathbf{x}_f^*$, starting at time $t_0 = 0$. The objective is to find the optimal control $\mathbf{v}^*(t)$ that minimizes total energy consumption and satisfies the boundary conditions:

$$\mathbf{v}^*(t) = \underset{\mathbf{v}(t)\in\mathcal{V}, t_f\in\mathbb{R}^+}{\arg\min} \int_{t_0}^{t_f} \dot{E}(t)\, dt \tag{6}$$

subject to:

$$\mathbf{x}(t_0) = \mathbf{x}_0^* \tag{7}$$

$$\mathbf{x}(t_f) = \mathbf{x}_f^* \tag{8}$$

and the dynamics and control constraints defined in (1) and (2). This formulation implicitly treats t_f (the final time) as an unknown, making it a free parameter also solved by the optimal control problem. When $\alpha = 0$, this formulation reduces to the minimum time problem.

III. Methods

To solve the optimal control problem, we apply PMP, which transcribes the infinite-dimensional optimization problem (6) into a TPBVP for a Hamiltonian system of coupled ODEs. Section III-A describes the application of PMP to the AUV trajectory optimization problem, deriving the optimal control policy, adjoint dynamics, and the associated TPBVP. Section III-B describes the numerical procedures used to solve the BVP in this work. Section III-C discusses limitations of the proposed numerical procedure. Additional background on PMP is presented in Appendix A. We also include Appendix B to highlight how PMP can be applied to related problems in AUV planning.

A. Theory

1) Application of PMP: To apply Pontryagin's Maximum Principle to the optimal control problem, we introduce the adjoint state vector (also referred to as the "co-state") $\mathbf{p}(t) : \mathbb{R} \mapsto \mathbb{R}^d$, which varies with time and has the same dimension as the state vector $\mathbf{x}$. The Hamiltonian is defined as a function of time, the state $\mathbf{x}$, the control $\mathbf{v}$, the adjoint state $\mathbf{p}$, and the energy consumption $\dot{E}$ (the running cost) as

$$H(t, \mathbf{x}, \mathbf{v}, \mathbf{p}) = \langle \dot{\mathbf{x}}, \mathbf{p} \rangle - \dot{E}, \tag{9}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors. Inserting the dynamics (1) and power model (4), the Hamiltonian of our problem is

$$H(t, \mathbf{x}, \mathbf{v}, \mathbf{p}) = \langle \mathbf{v}, \mathbf{p} \rangle + \langle \mathbf{u}, \mathbf{p} \rangle - \alpha \|\mathbf{v}\|^3 - \beta. \tag{10}$$

The *canonical equations*, describing the dynamics of the state and adjoint state, are given in terms of the Jacobian of the Hamiltonian:

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}} = \mathbf{v} + \mathbf{u}, \tag{11}$$

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}} = -\nabla_{\mathbf{x}} \mathbf{u}^\top \mathbf{p}, \tag{12}$$

where $\nabla_{\mathbf{x}}$ denotes the spatial gradient. Note that (11) is consistent with the original dynamics (1).

PMP asserts that, as a necessary condition, the optimal control $\mathbf{v}^*$ maximizes the Hamiltonian (10) over the permissible control set at all points along the trajectory. We express the control vector $\mathbf{v}$ in terms of its magnitude $\|\mathbf{v}\|$ and its direction $\hat{\mathbf{v}}$ to obtain

$$H(t, \mathbf{x}, \|\mathbf{v}\|, \hat{\mathbf{v}}, \mathbf{p}) = \|\mathbf{v}\|\langle \hat{\mathbf{v}}, \mathbf{p}\rangle + \langle \mathbf{u}, \mathbf{p}\rangle - \alpha\|\mathbf{v}\|^3 - \beta. \tag{13}$$

Note that the control direction $\hat{\mathbf{v}}$ appears only in the first term $\langle \hat{\mathbf{v}}, \mathbf{p}\rangle$. From the Cauchy-Schwarz inequality, this inner product is maximized when the two vectors $\mathbf{v}$ and $\mathbf{p}$ are aligned. Since there is no constraint on the vehicle orientation, the direction of the optimal control is therefore aligned with the adjoint vector at all points along the optimal trajectory: $\hat{\mathbf{v}}^* = \hat{\mathbf{p}}$. The optimal control magnitude $\|\mathbf{v}\|$ then maximizes

$$H(t, \mathbf{x}, \|\mathbf{v}\|, \mathbf{p}) = \|\mathbf{v}\|\|\mathbf{p}\| + \langle \mathbf{u}, \mathbf{p}\rangle - \alpha\|\mathbf{v}\|^3 - \beta. \tag{14}$$

To maximize (14) with respect to the magnitude $\|\mathbf{v}\|$, we set the partial derivative to zero:

$$\frac{\partial H}{\partial \|\mathbf{v}\|} = \|\mathbf{p}\| - 3\alpha\|\mathbf{v}\|^2 = 0, \tag{15}$$

so that, if feasible, the magnitude of the optimal control is

$$\|\mathbf{v}^*\| = \sqrt{\frac{\|\mathbf{p}\|}{3\alpha}}. \tag{16}$$

Considering the constraint (2), the optimal control as a function of time is:

$$\mathbf{v}^* = \begin{cases} \sqrt{\frac{\|\mathbf{p}\|}{3\alpha}}\hat{\mathbf{p}}, & v_{\min} \le \sqrt{\frac{\|\mathbf{p}\|}{3\alpha}} \le v_{\max} \\ v_{\min}\hat{\mathbf{p}}, & \sqrt{\frac{\|\mathbf{p}\|}{3\alpha}} < v_{\min} \\ v_{\max}\hat{\mathbf{p}}, & \sqrt{\frac{\|\mathbf{p}\|}{3\alpha}} > v_{\max} \end{cases}. \tag{17}$$

Finally, since (9) explicitly depends on time and t_f is free, the following transversality condition must also be satisfied:

$$H(t_f, \mathbf{x}_f, \mathbf{p}_f) = 0. \tag{18}$$

In the case of a static flow field, where $\mathbf{u}(t, \mathbf{x}) \equiv \mathbf{u}(\mathbf{x})$, the Hamiltonian becomes time-independent. Thus, it can be shown that $H = 0$ for all t, and (18) can be evaluated at any state along the trajectory [10].

The Hamiltonian BVP is defined by the dynamics (11) and (12), the optimal control (17), the initial and final vehicle positions (7) and (8), and the transversality condition (18). PMP establishes a first-order necessary condition for optimality; therefore, trajectories that satisfy this BVP are candidates for globally optimal solutions to the path-planning problem.

2) Interpreting the Optimal Control Policy: It is notable that β does not directly influence the optimal control in (17). Intuitively, we expect that with a higher hotel load, the optimal control will correspond to a faster speed to minimize travel time. The effect of β becomes more apparent through (18). Assuming the control is unconstrained, substitute (17) into (10) and evaluate the transversality condition:

$$\langle \mathbf{v}^*, \mathbf{p}\rangle + \langle \mathbf{u}, \mathbf{p}\rangle - \alpha\|\mathbf{v}^*\|^3 - \beta = 0 \tag{19}$$

$$\langle \mathbf{v}^*, 3\alpha\|\mathbf{v}^*\|^2\hat{\mathbf{v}}^*\rangle + \langle \mathbf{u}, 3\alpha\|\mathbf{v}^*\|^2\hat{\mathbf{v}}^*\rangle - \alpha\|\mathbf{v}^*\|^3 - \beta = 0 \tag{20}$$

$$2\alpha\|\mathbf{v}^*\|^3 + 3\alpha u_\theta\|\mathbf{v}^*\|^2 - \beta = 0 \tag{21}$$

where $u_\theta := \langle \mathbf{u}, \hat{\mathbf{v}}\rangle$ is the along-track projection of the current. Note that u_θ aligns with the vehicle's heading, not its course in the global reference frame. From (21), we see that $\|\mathbf{v}^*\|$ must increase for larger β, as expected.

In the absence of currents ($\mathbf{u} = \mathbf{0}$), the optimal speed is determined as:

$$\|\mathbf{v}_0\| = \left(\frac{\beta}{2\alpha}\right)^{1/3}. \tag{22}$$

Normalizing (21) by $\|\mathbf{v}_0\|$, we obtain the cubic equation:

$$\bar{v}^{*3} + \frac{3}{2}\bar{u}_\theta\bar{v}^{*2} - 1 = 0 \tag{23}$$

where $\bar{v}^*$ is the normalized optimal speed and $\bar{u}_\theta$ is normalized along-track current speed. This normalization allows us to separate the effects of the vehicle parameters, α and β, from the effects of

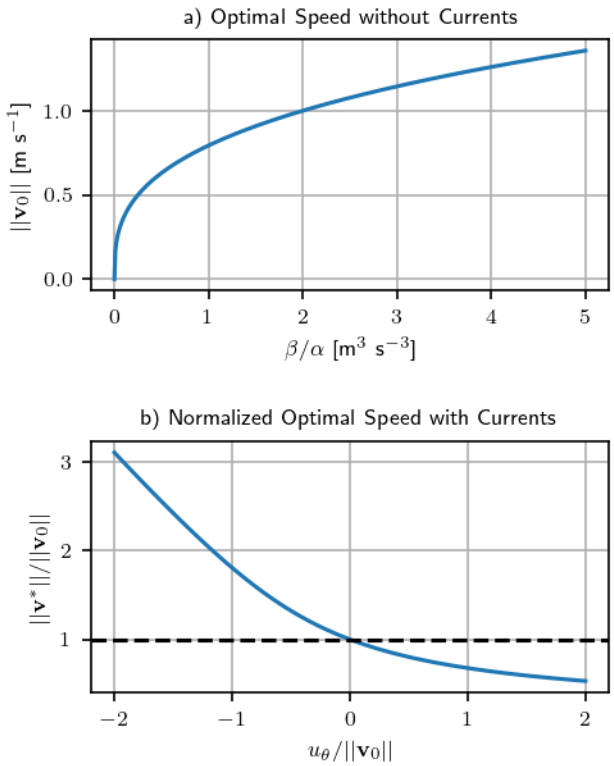

Fig. 1. a: No-current optimal speed as a function of β/α, the ratio of hotel load to the dynamic power coefficient. b: Normalized optimal speed through water as a function of normalized along-track current speed. Both values are normalized by the no-current optimal speed.

the currents. The ratio β/α determines the nominal speed of the vehicle (the optimal speed in no currents), shown in Fig. 1a. Fig. 1b then shows how the current's along-track component increases the optimal speed when positively aligned and decreases it when opposed to the vehicle's orientation. The cross-component of the current does not affect the optimal speed. Because the transversality condition is applied only at the final time boundary, the vehicle can operate at speeds different from the locally optimal value. This flexibility allows the vehicle to exploit transient features in the flow field. For example, when $H > 0$ the instantaneous cost rate is greater than what is locally required and the vehicle will appear to go faster than expected.

B. Numerical Implementation

To solve the BVP described in Section III-A, we propose a straightforward numerical framework based on pure shooting [46]. This approach follows naturally from the transcription provided by PMP and involves two standard numerical procedures: 1) a 4th order Runge-Kutta method to integrate the canonical equations (11) and (12), and 2) the Newton-Raphson method for root-finding (to satisfy the boundary conditions).

1) Transcription: To solve the stated TPBVP via shooting, one needs to determine $\mathbf{p}_0$ and t_f ($d+1$ unknowns) such that integration of the canonical equations satisfies the boundary conditions. This process can be expressed as finding zeros of the shooting function S_1:

$$S_1 : \begin{pmatrix} t_f \\ \mathbf{p}(t_0) \end{pmatrix} \mapsto \begin{pmatrix} H(t_f, \mathbf{x}(t_f), \mathbf{p}(t_f)) \\ \mathbf{x}(t_f) - \mathbf{x}_f^* \end{pmatrix}. \tag{24}$$

A common challenge in shooting methods is the determination of a suitable initial guess for the adjoint state vector, $\mathbf{p}(t_0)$. Previous studies have addressed this issue with various approaches. For example, Dobrokhodov et al. [27] applied a homotopy method by scaling the wind field for aerial vehicle applications. Cristiani and Martinon [11] suggested a general approach for initializing PMP solutions using a low-resolution solution to the HJB equation.

In this application, we will initialize the adjoint vector using the transversality condition. Consequently, we can reduce the search space to d dimensions which are physically intuitive. Recall that for the Hamiltonian (10), the optimal control and the adjoint state vector are co-linear. Specifically, for an optimal trajectory $\hat{\mathbf{v}}^*(t_f) = \hat{\mathbf{p}}(t_f)$. Then, the magnitude $\|\mathbf{p}(t_f)\|$ can be found by substituting (16) into (21), yielding:

$$\frac{2}{3}\sqrt{\frac{1}{3\alpha}}\|\mathbf{p}(t_f)\|^{3/2} \\ + u_\theta(t_f, \mathbf{x}_f^*, \hat{\mathbf{v}}^*(t_f))\|\mathbf{p}(t_f)\| - \beta = 0. \tag{25}$$

Therefore the final vehicle orientation, which is determined by $d-1$ angles, completely determines the final adjoint state vector $\mathbf{p}(t_f)$ for optimal trajectories. The updated transcription involves finding the zeros of a refined shooting function S_2:

$$S_2 : \begin{pmatrix} t_f \\ \hat{\mathbf{v}}(t_f) \end{pmatrix} \mapsto \left(\mathbf{x}(t_0) - \mathbf{x}_0^*\right) \tag{26}$$

where the adjoint vector is initialized with $\hat{\mathbf{v}}(t_f)$, as the solution to (25). The canonical equations are integrated backwards in time from t_f to 0. For convenience, we define $\mathbf{y} := (t_f, \hat{\mathbf{v}}(t_f)) = (t_f, \theta_f, \phi_f)$ as the input to the shooting function S_2.

2) Solution Procedure: To find both local and global solutions to the trajectory planning problem, we use Algorithm 1, based on the transcription described by (26). Although the algorithm is specific to an environment with a 3-D, time-varying current field, it is straightforward to adapt to lower dimensions. To initialize the algorithm we require a regular grid of inputs to the shooting function $\mathbf{Y} \subseteq \mathbb{R} \times [0, 2\pi) \times [-\pi/2, \pi/2]$, a distance tolerance r used to identify candidates from $\mathbf{Y}$ for subsequent refinement, and the parameters of the OCP: $\mathbf{x}_0^*$, $\mathbf{x}_f^*$, $\mathbf{u}$, α, and β. The algorithm returns a set, $\gamma^\dagger = \{\mathbf{y} | S_2(\mathbf{y}) = \mathbf{0}\}$, containing zeros of the shooting function (local minima), as well as the vector $\mathbf{y}^* \in \gamma^\dagger$ corresponding to final state of the *global* minimum energy trajectory.

Algorithm 1 Solve BVP: 3-D, time-varying

Input: Θ, Φ, $\mathcal{T}_f$, r, ($\mathbf{x}_0$, $\mathbf{x}_f$, $\mathbf{u}$, α, β, d)

% Shoot exhaustive extremal trajectories
for each $\mathbf{y}_{ijk}$ in $\Theta \times \Phi \times \mathcal{T}_f$ **do**
  d_{ijk}, e_{ijk} = $\texttt{shooting}(\mathbf{y}_{ijk})$
% Find local minima of d within tolerance r
$\gamma \leftarrow ()$
for each $(i, j, k), (\theta, \phi, t_f)$ in $\Theta \times \Phi \times \mathcal{T}_f$ **do**
  if d_{ijk} is a local minimum **and** $d_{ijk} \leq r$ **then**
    γ.append((θ, ϕ, t_f))
% Perform root-finding and check convergence
$\gamma^\dagger \leftarrow ()$
$\bar{\mathbf{e}} \leftarrow ()$
for each $\mathbf{y}$ in γ **do**
  $\bar{\mathbf{y}}$, $\bar{e}$, did_converge = $\texttt{newt}(\mathbf{y})$
  $\gamma^\dagger$.append((θ, ϕ, t_f))
  $\bar{\mathbf{e}}$.append($\bar{e}$)
% identify global vs local minima and return
$\mathbf{y}^* = \arg\min_{\mathbf{y} \in \gamma^\dagger} \bar{\mathbf{e}}$
return $\mathbf{y}^*, \gamma^\dagger$

The algorithm relies on two auxiliary functions. The first, $(d, e) = \texttt{shooting}(\mathbf{y})$, implements the shooting function S_2 (26). It integrates a proposed final state backward to t_0 using a 4th order Runge-Kutta integrator and returns the euclidean distance error, $d := \|S_2\|_2$, at the initial boundary condition, along with the total energy consumed along the trajectory, e. For cases where the current field is gridded, spatial gradients of the current are precomputed using finite differences and queried during integration via linear interpolation.

The second auxiliary function, $(\bar{\mathbf{y}}, \bar{e}, \text{flag}) = \texttt{newt}(\mathbf{y})$, uses a Newton-Raphson method to find zeros of the shooting function. It returns updated values of the input, $\bar{\mathbf{y}}$, along with the corresponding energy consumption $\bar{e}$, and a convergence flag indicating whether the function has met the desired numerical precision. The gradients of the shooting function, $\frac{\partial S_2}{\partial \mathbf{y}}$, are computed using finite differences.

The solution procedure begins by integrating the canonical equations across a range of final conditions to identify reasonable initial guesses for root-finding. The user specifies discrete regular grids for heading angles Θ, pitch angles Φ, and final times $\mathcal{T}_f$, which together define the grid of initial guesses, $\mathbf{Y} = \mathcal{T}_f \times \Theta \times \Phi$. Each element of this grid is then evaluated using the shooting function, and the Euclidean distance error at the initial boundary is stored in a grid, D of the same dimensions as $\mathbf{Y}$.

Next, the local minima of D are identified by comparing adjacent grid points. Local minima within the specified tolerance r are then used to initialize $\texttt{newt}$, which attempts to converge to a zero of the shooting function. Once the set of initial conditions that satisfy the boundary conditions is found, the corresponding trajectory, control policy, and energy consumption can be determined. The trajectory that consumes the least energy consumption is selected as the globally optimal solution.

C. Considerations for Alternate Numerical Frameworks

In the following sections, Algorithm 1 is shown to be sufficient for solving problems of varying levels of complexity. Section IV demonstrates its efficacy for trajectories in various analytically defined flows and that it compares favorably in run-time to a PDE-based level-set method for a 3-D, time-varying current field. In section V, the algorithm successfully computes optimal trajectories through realistic current fields from high-resolution ocean models. In these flows, the algorithm will converge to the globally-optimal solution with a relatively coarse (1° or 2°) discretization of trajectory angles. The root-finding procedure requires on the order of 100 iterations total, with order 10 iterations used to converge from each initial guess that is within the prescribed tolerance.

However, it is well-known that simple shooting can demonstrate an extreme sensitivity to initial conditions that precludes its application to certain problems [46]. In particular, the authors found that Algorithm 1 can break down for 3-D ocean currents with strong vertical gradients, where the vertical length scale is much shorter than the horizontal length scale and the trajectory length exceeds both environmental length scales. For such problems, a variety of numerical techniques such as multiple shooting [47] and collocation methods [48], [49] have been developed to overcome the instability of simple shooting for solving TPBVPs. Note that the initialization step for multiple shooting or collocation will require an initial guess for the adjoint vector at multiple points along the trajectory. Therefore, the simplified shooting function (26) is not applicable, as it is a result of applying the Hamiltonian boundary condition. Given the simplicity and speed of Algorithm 1, and its success on the examples presented here, it is thought to be an attractive alternative to existing, more computationally demanding PDE-based methods for many practical optimal trajectory calculations. The use of alternative TPBVP solvers to solve for optimal trajectories in more challenging environments is left for future work.

IV. Canonical Examples

Our methodology is illustrated through examples of time-optimal trajectory planning in stationary flow fields. Solutions to a 2-D, 3-D, and time-varying, 3-D problem are shown in Fig. 2, Fig. 3, and Fig. 4, respectively. The governing equations for time- and energy-optimal trajectories are nearly equivalent (see Appendix B-B): adjoint dynamics (12) remain the same, and the optimal direction is co-linear with the adjoint state vector, but the vehicle speed is fixed instead of being a function of $\|\mathbf{p}\|$.

A. 2-D, stationary currents

In the first example, we examine solutions in a static, 2-D current field,

$$\mathbf{u}(x, y) = [A\sin(\pi y),\, 0], \tag{27}$$

which was previously studied by Davis et al. [25]. In this example, the boundary conditions are $\mathbf{x}_0^* = (0,0)$ and $\mathbf{x}_f^* = (0,1)$, and the vehicle travels at a fixed speed of $\|\mathbf{v}\| = 1.0$. We solve this problem using Algorithm 1, which is initialized by the grid Θ, with 100 evenly spaced heading angles between $[0, 2\pi)$, and the final time grid $\mathcal{T}_f$, which ranges from $[0, 2.5]$ with a time step of 1×10^{-3}. As the vehicle lies in the x-y plane in this example, there is no search over the pitch angle Φ.

Fig. 2 shows locally time-optimal trajectories in the current field (27) for three current strength values $A = [0.0, 0.75, 1.5]$. In the top row, gray lines correspond to optimal trajectories initialized by each value of Θ. Since the current field is stationary, for a given final heading, trajectories for all values of t_f follow the same path. This is equivalent to integrating the canonical equations backwards in time from $t = 0$ to $t = -2.5$. Therefore, increasing the value of t_f has the effect of moving along a given trajectory away from the end point. The trajectories that intersect the starting point (solutions to the TPBVP) are shown in color. The lower row contours the distance of the locally optimal trajectories shown in the upper row from the desired starting point $(x_0^*, y_f^*) = (0, 0)$ as a function of time along the trajectory (or final time) and final vehicle orientation.

When $A = 0.0$, the trajectories are straight lines: in the absence of currents, time-optimal paths are shortest-distance paths. As A increases, the trajectories become distorted due to the influence of stronger currents. For $A = 0.75$, a single globally optimal trajectory curves slightly against the current. On the space of final conditions, it is clear that there is still only one solution, although the contours of D become more complicated. Additionally, there is an extra zero-contour in the y_0 boundary condition, which corresponds to the trajectories that are bending downward in the top half of the figure. When $A = 1.5$, these features are more pronounced, leading to multiple zeros of the shooting function, as seen in Fig. 2f. The central section of all trajectories are nearly identical, with the strong currents advecting the vehicle from left to right. However, the solutions differ near the boundary conditions. Both the green and the blue solutions exploit the strong negative current outside the boundary conditions in the y direction, allowing the vehicle to drift quickly in the opposite direction of the central segment. The controls for the green and blue trajectories are rotationally symmetric, and their travel times are equivalent due to the symmetry of the current field with respect to an exchange of the initial point and end point. Although their trajectories are longer in

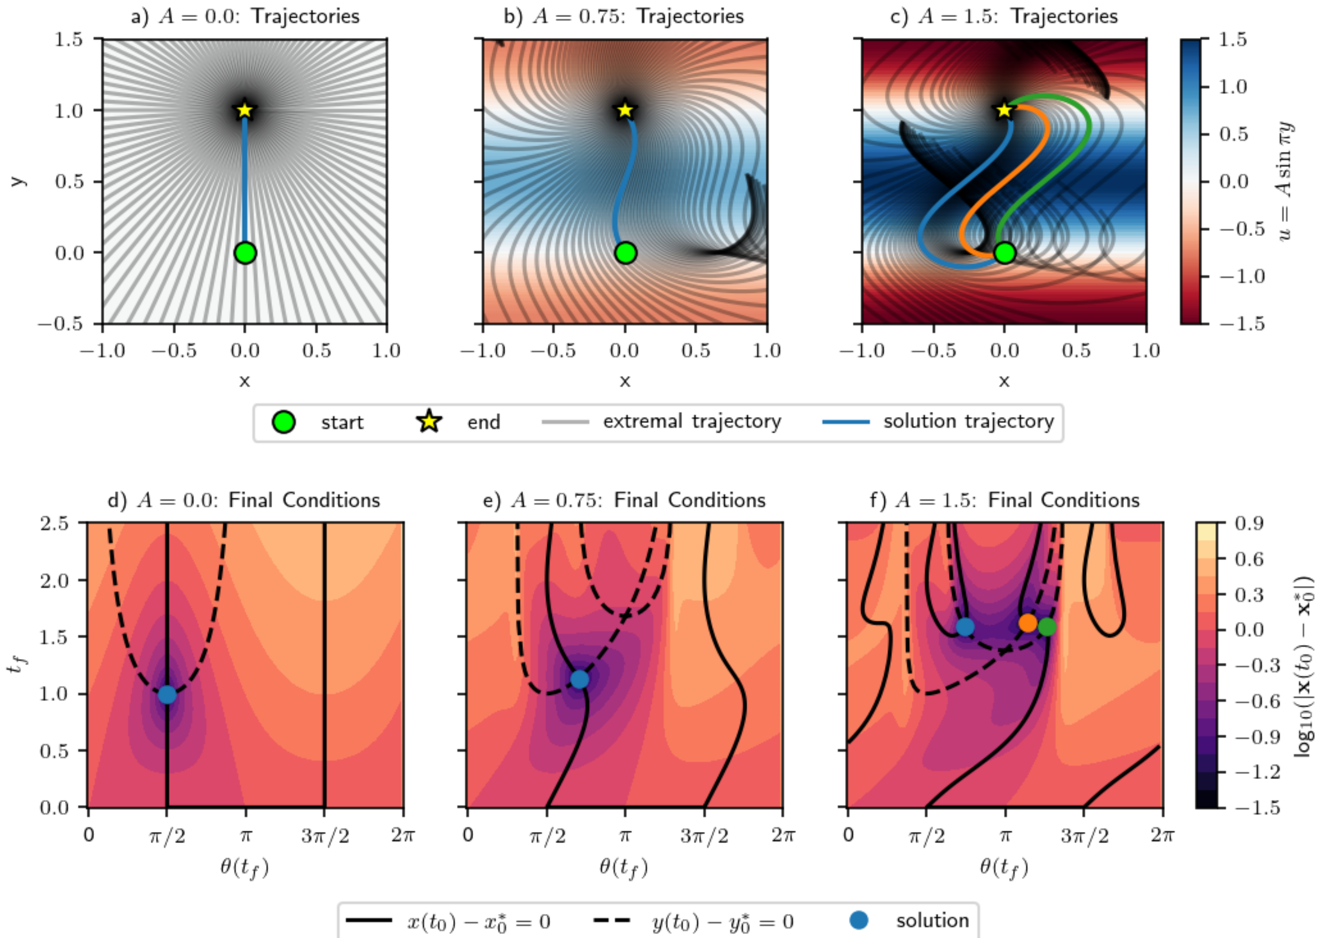

Fig. 2. a-c: Minimal time, extremal trajectories (gray) and locally-optimal trajectories (color) are shown in a horizontal current field, $u = A\sin(\pi y)$ and $A = [0.0, 0.75, 1.5]$, for a vehicle with constant speed $\|\mathbf{v}_0\| = 1.0$. The strength of the horizontal current is shown in color. Note that three locally optimal solutions exist for $A = 1.5$. d-f: Contours of the euclidean distance error with respect to the $\mathbf{x}_0^*$ boundary condition are shown on the space of final conditions (heading and time). Zero-contours of the x_f and y_f boundary conditions are overlaid with solid and dashed black lines, respectively.

distance, they are faster than the orange solution, which follows a more direct path.

B. 3-D, stationary currents

For the second example, we demonstrate solutions for time-optimal trajectories in a 3-D, stationary current field. For this example, we augment the previous current field (27) to include a dependence on depth:

$$\mathbf{u}(x, y, z) = \left[A\sin(\pi y)\exp\left(\frac{-z^2}{2\ell^2}\right),\ 0,\ 0 \right]. \quad (28)$$

The parameter ℓ, is a vertical length scale such that the magnitude of the flow decreases symmetrically out of the plane $z = 0$. The boundary conditions are the same as before with the addition of a depth constraint $z_0^* = z_f^* = 0$. Algorithm 1 uses the same grids for Θ and $\mathcal{T}_f$ as in the 2-D case. The grid Φ, is initialized with 200 evenly spaced pitch angles between $[-\pi/2, \pi/2]$. We solve this problem in three different current fields where $A = 1.5$ is fixed, and $\ell = [1.0, 0.4, 0.25]$, which represents stronger vertical shear in the flow.

Results for these three current fields are shown in Fig. 3. When $\ell = 1.0$, and the currents are effectively uniform in depth relative to the length of the transit, the optimal trajectories are equivalent to those shown in Fig. 2c. As we increase ℓ, there are additional extremal trajectories that navigate out of the plane to avoid the strong cross-current. As in the planar example, due to symmetry in the current field, the trajectories that explore $z > 0$ and $z < 0$, respectively, have the same final time and distance traveled. When $\ell = 0.4$, these trajectories are only

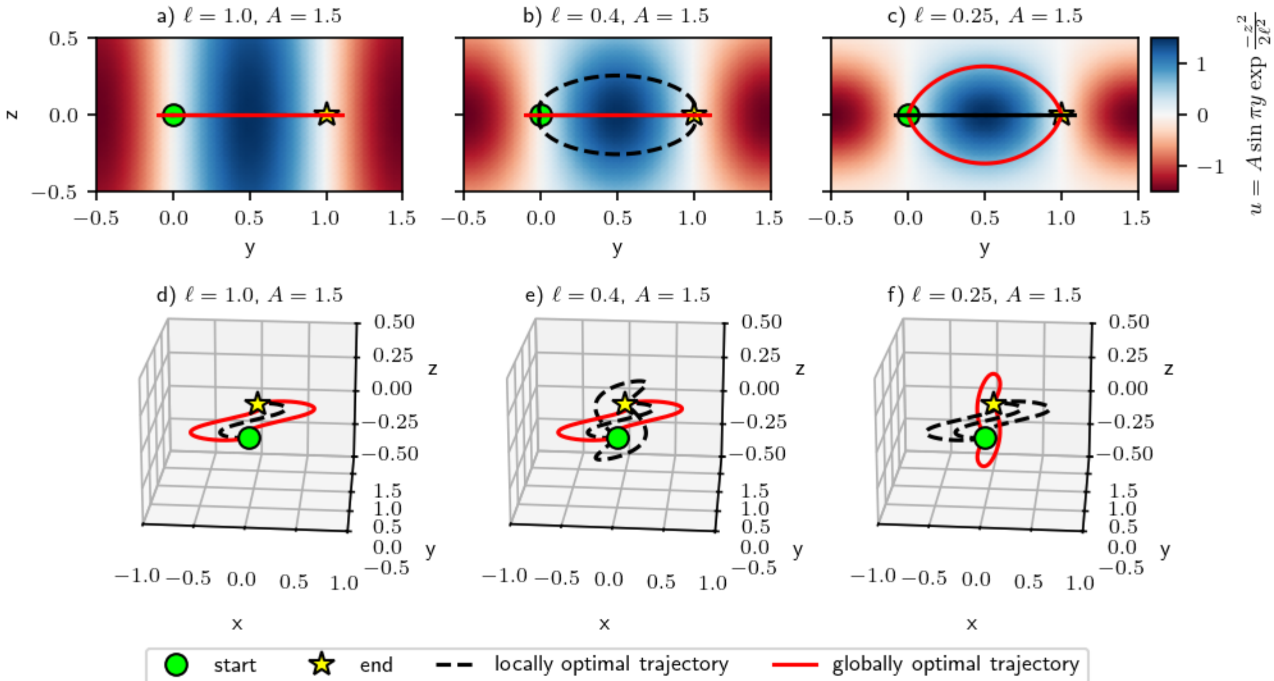

Fig. 3. a-c: The horizontal current field $u = A\sin(\pi y)\exp\left(-z^2/2\ell^2\right)$ is shown as a function of y and z for three different length scales $\ell = [1.0, 0.4, 0.25]$. 3-D, time-optimal trajectories are calculated between $\mathbf{x}_0^* = (0,0,0)$ and $\mathbf{x}_f^* = (0,1,0)$ and projected onto the y-z plane. Once ℓ is sufficiently small (b-c) there are optimal trajectories that deviate from the plane $z = 0$. d-f: The optimal trajectories corresponding to the same current field in the row above, shown in full 3-D. Locally optimal trajectories are shown by dashed lines, while the globally optimal (multiple due to symmetry) are shown in red.

locally optimal when compared to the outermost transects on the plane $z = 0$. The currents are still relatively strong once the vehicle has moved vertically, which is reflected in the trajectories which deviate in the x-y plane as before. In the most extreme case, $\ell = 0.25$, it becomes globally optimal to deviate from $z = 0$. In these trajectories, navigation away from $x = 0$ is minimal, and time-savings are achieved by moving vertically away from the strong currents.

C. 3-D, unsteady currents

As a final example, we analyze trajectories through a 3-D, time-varying current field which was previously considered by Kulkarni and Lermusiaux [17]. This provides a comparison of our method to a PDE-based, level-set method with reported run-times from [17]. The unsteady flow is defined by

$$u_x = -\pi A \sin(\pi f(x))\cos(\pi y) \tag{29}$$

$$u_y = \pi A \cos(\pi f(x))\sin(\pi y)\frac{df}{dx} \tag{30}$$

$$u_z = xy(2-x)(1-y)(1+\epsilon\sin(\omega t)-x), \tag{31}$$

where

$$f(x) = a(t)x^2 + b(t)x \tag{32}$$

$$a(t) = \epsilon \sin(\omega t) \tag{33}$$

$$b(t) = 1 - 2\epsilon\sin(\omega t). \tag{34}$$

The flow field does not depend on z. However, there is a non-zero vertical component (31) which causes trajectories to deviate from the x-y plane. The horizontal components of the flow, shown in Fig. 4d, approximate a geostrophic double gyre. The vertical component, shown in Fig. 4e, is positive in one gyre and negative in the other. The gyre system oscillates back and forth across $x = 1$ with an amplitude of ϵ and a frequency ω. For this example, the parameters of the flow are $A = 0.1$, $\epsilon = 0.1$, and

TABLE I
OPTIMAL TRAJECTORY TIME (t_f), AND COMPUTATION TIMES REQUIRED FOR ALL-TO-ONE BROADCAST THROUGH 3-D, TIME-VARYING DOUBLE GYRE CURRENTS.

Point	**Location** (x, y, z)	**Optimal** t_f	Initial Comp. Time [s]	Convergence Comp. Time [s]
goal	$(0.65, 0.35, 0.75)$	–	4.59	–
start 1	$(0.85, 0.75, 0.6)$	8.51	–	0.17
start 3	$(1.75, 0.30, 0.40)$	15.62	–	0.92
start 2	$(1.30, 0.40, 0.50)$	15.84	–	0.41

$\omega = \pi/5$, in agreement with [17].

We compute time-optimal trajectories for a vehicle with a fixed speed $\|\mathbf{v}\| = 0.05$, which is extremely under-actuated in this environment. The maximum horizontal current speeds exceed 0.3, and maximum vertical speeds exceed 0.1. We consider three vehicles traveling to a common goal point starting from three different starting locations at $t = 0.0$. The start locations, optimal trajectory transit time t_f, and run time are summarized in Table I. The optimal trajectories are shown in Fig. 4a-c and clearly follow the general circulation of the flow dominated by the structure of the double gyre. In the vertical direction, notice that the vehicles starting in the counter-clockwise rotating gyre descend until approximately $x = 0.9$ and then increase their depth to the goal. This is consistent with the direction of the vertical component of the current, and the midpoint between the two gyres which has shifted with time.

In [17], three trajectories were computed from a single starting point to three goal points, whereas in this work three trajectories are computed for three starting points and a single shared goal point. This discrepancy reflects a fundamental difference between level-set methods and the TPBVP method presented here. Algorithm 1 naturally lends itself to computing trajectories to a single goal point from multiple start points (an "all-to-one" broadcast) because the shooting function (26) is initialized *at the goal point*. Therefore, the same initial grid of extremal trajectories (Algorithm 1 1-3) can be reused to initialize the root-refinement process for many trajectories. In contrast, the level-set method naturally lends itself to the reciprocal problem: computing trajectories from one common starting point to several goal points (a "one-to-all" broadcast"). This is because level-sets propagate forward in time away from the starting point. Then, an inexpensive backtracing step can compute the optimal trajectory to any selected goal point within the region explored by the level sets. This gives level-set methods a computational advantage for one-to-all broadcasts, whereas the PMP approach taken here offers a computational advantage for "all-to-one" broadcasts. We note that when the current field is stationary, and the transversality condition can be evaluated at the start point, our PMP-based approach can solve the broadcast problem in either direction equivalently using Algorithm 1.

To compare the computation time of Algorithm 1 to the stated runtime in [17][2], optimal trajectories are computed for the same current field (29)–(31). To put the two methods on an equal footing, the goal and start locations used in [17] are exchanged to convert the problem from a one-to-all broadcast to an all-to-one broadcast. We initialize Algorithm 1 with a uniform grid over $\mathbb{S}^2$ with a discretization of 0.342 radians in heading and pitch angles, respectively. We use a grid of final times $\mathcal{T}_f$ between $[5, 20]$ with a time step of 1.0. The initial shooting is computed in 4.59 seconds, and root-finding for each start point took less than 1.0 second respectively. In comparison, a level-set expansion for the same current field took 82.33 seconds [17]. Because our proposed method is not guaranteed to produce an optimal solution, whereas PDE approaches are, a formal comparison of algorithmic complexity is not possible. As a final note, we point out that the method presented here also lends itself to parallelization. The most expensive aspect of our algorithm is the shooting of locally

[2] Our computations were performed on a Macbook Pro with an Apple M2 Max, 12-core processor. Algorithm 1. was implemented in Python using Numba [50] for just-in-time compilation. The computations from [17] were completed on a computer with a single quad-core Intel i7-4790 CPU at 3.60 GHz. The level-set method was implemented in MATLAB with internal vectorization/parallelization of matrix-vector operations enabled.

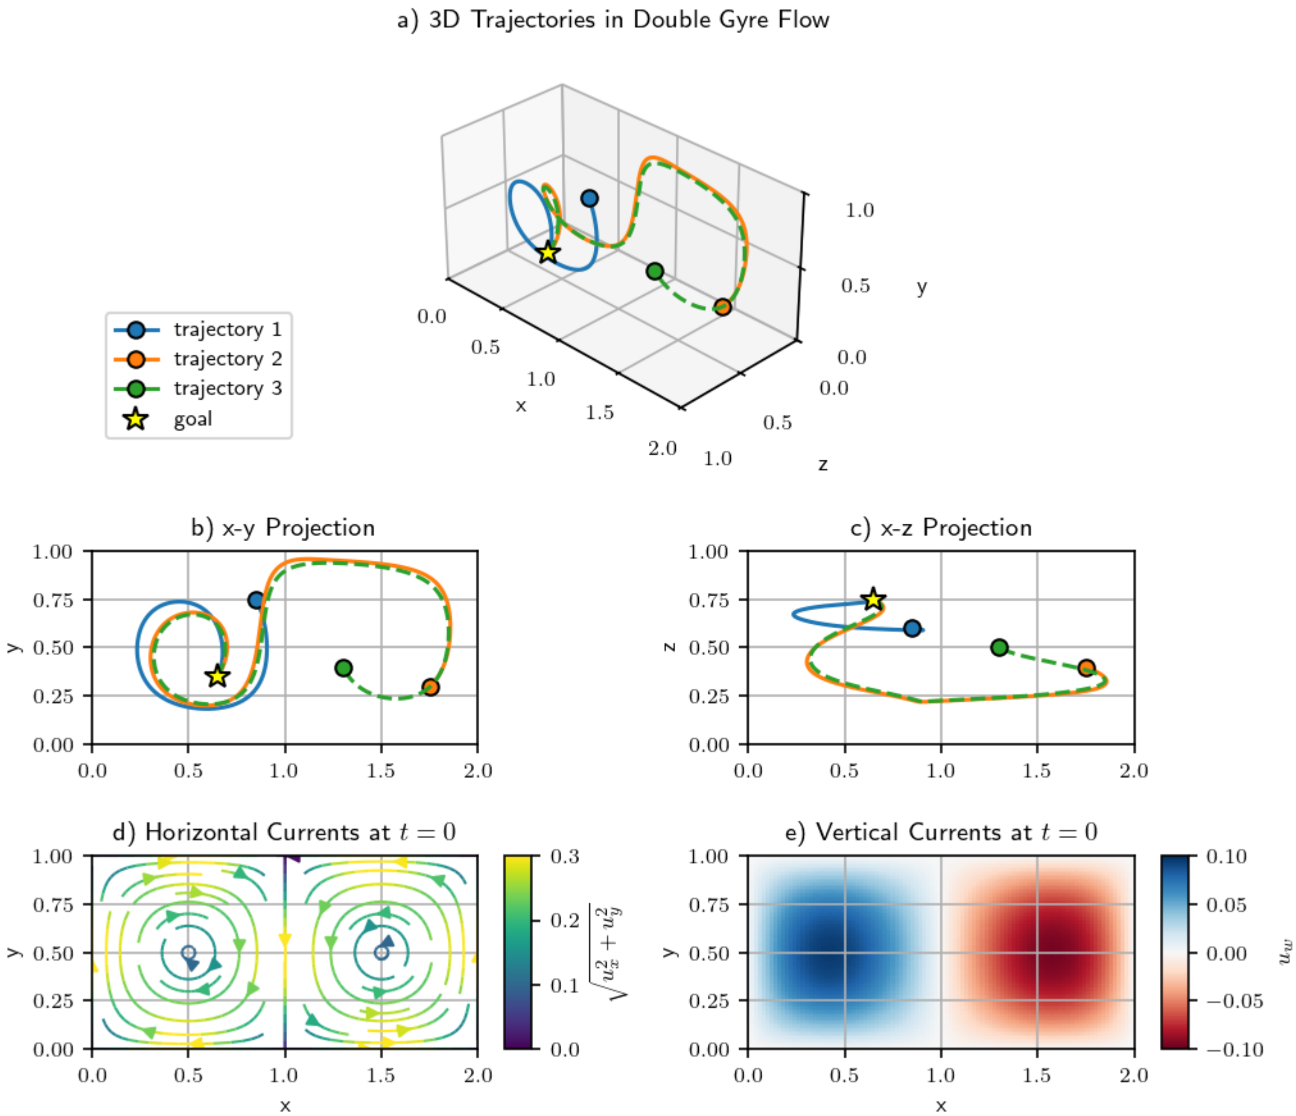

Fig. 4. a-c: Time-optimal trajectories for vehicles starting at $t = 0$ from three different starting points to a common goal (yellow star) viewed in 3-D (a), projected onto the x-y plane (b), and onto the x-z plane (c).d-e: The horizontal components (d) and vertical component (e) of the current field at time $t = 0$ as a function of x and y.

optimal trajectories for each value in the initial grid. However, since each trajectory is independent from one another, this step could easily be distributed across processing units. The initial trajectories can then be gathered for use in the root-finding step to compute an optimal trajectory for a particular starting point.

V. Simulated Experiments

In this section, we present two case studies using ocean currents derived from high-resolution regional ocean simulations. In both cases, we consider time-dependent surface currents (2-D) which are relevant for an AUV operating near the surface and sampling the ocean at a fixed depth. These results can be generalized to surface vehicles using an appropriate energy consumption model.

For each case study, we compare the behavior of different vehicles represented by a range of no-current optimal speeds, $\|\mathbf{v}_0\|$. We present results for energy-optimal and time-optimal trajectories, as well as a baseline planning strategy in which the vehicle follows a straight line, adjusting its heading while maintaining a constant speed through water. The time-optimal case is equivalent to energy-optimal control for a vehicle with a fixed speed

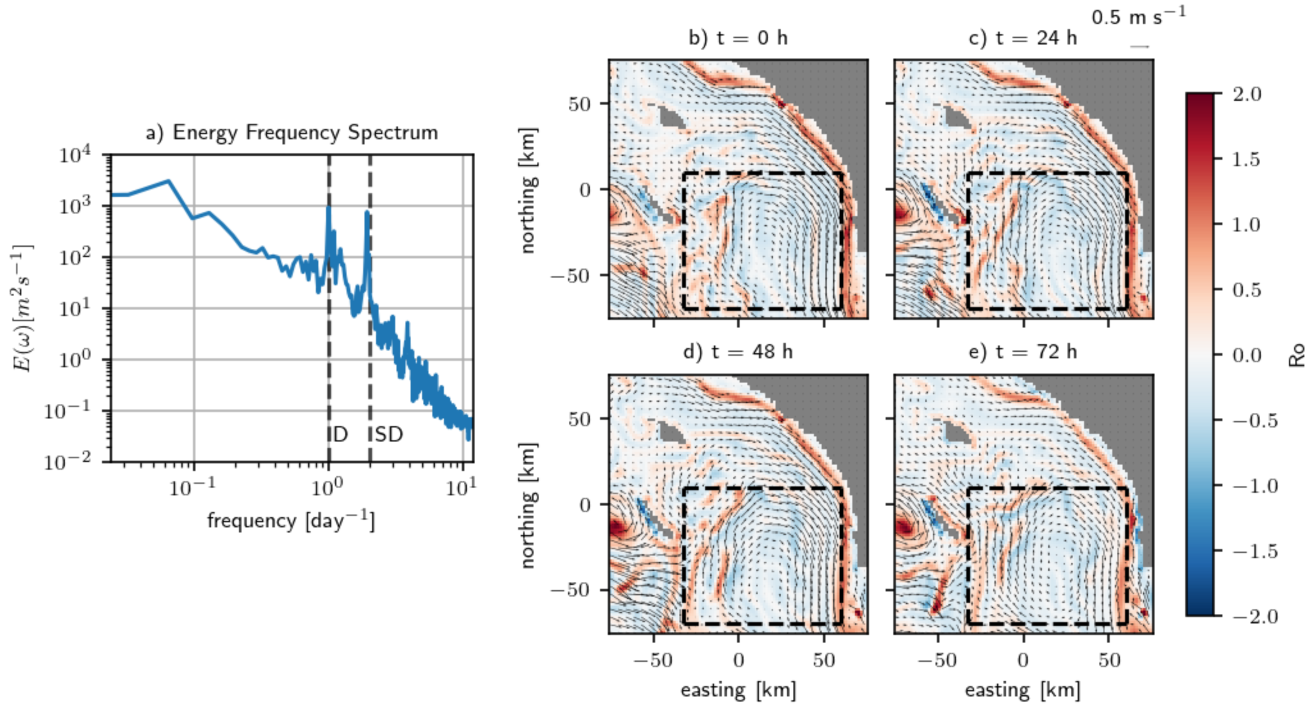

Fig. 5. a: Spectral characteristics of surface currents in the Southern California Bight. Kinetic energy frequency spectra are computed at each grid cell of the boxed region and averaged in space. Vertical lines annotate diurnal (D), and semi-diurnal (SD) frequencies, respectively. b-e: The temporal evolution of currents in the Southern California Bight over 72 hours. Maps colored by Rossby number (Ro) are overlaid with quiver arrows for surface currents. Positive and negative Ro represent cyclonic and anti-cyclonic circulation, respectively.

equal to $\|\mathbf{v}_0\|$. For details on the time-optimal methodology refer to Appendix B-A, and for details on the implementation of the line-following policy, refer to Young et al. [51].

A. *Southern California Bight*

The first case study focuses on a long-duration transit in the Southern California Bight, from San Clemente Island to San Diego, CA. This transit is modeled after repeat, long duration transects that are completed by autonomous gliders for coastal monitoring over climatological time scales [52]. The boundary conditions are set in a local coordinate system with the origin at 33° 0.43' N, 117° 59.90' W, with starting position $\mathbf{x}_0^* = (-32.5, -20.0)$ km and final position $\mathbf{x}_f^* = (60.0, -40.0)$ km. Vehicle parameters α and β are chosen as pairs such that the no-current optimal speed through water covers the range $\|\mathbf{v}_0\| = [0.5, 2.0]$ m/s. To initialize Algorithm 1 we use a grid of initial conditions where $\mathcal{T}_f$ spans $[1\text{h}, 3\text{d}]$ with 0.5 h spacing and Θ_f spans $[-90°, 90°]$ with $1°$ spacing. The integrator used a 10-minute time step.

The ocean current field is obtained from a regional MITgcm [53] simulation with hourly temporal resolution and approximately 1.8-km horizontal resolution [54]. The currents in this region are dominated by mesoscale features while tidal contributions at diurnal (D) and semi-diurnal (SD) frequencies are significant but not as energetic, as shown in Fig. 5a. For a detailed discussion of the environmental considerations in this region for path planning, see Young et al. [55]. This experiment focuses on a 72-hour period with a persistent mesoscale eddy, superimposed with tidal variability, Fig. 5b-e. These features are highlighted by the Rossby number (Ro), defined as $\text{Ro} = \frac{\zeta}{f}$, where ζ, relative vorticity, is normalized by f, planetary vorticity which is a function of latitude.

In the case of no currents, when the optimal trajectory follows a straight line, both the time and energy consumption are deterministic, following the red dashed curve in Fig. 6a. As expected, increasing vehicle speed results in a shorter transit time but higher energy consumption. These results are normalized such that when $\|\mathbf{v}_0\| = 1.0$ m/s, the vehicle completes the trajectory in unit time and energy. The curve for energy-optimal trajectories, following colored circles in Fig. 6a, shows that slightly more energy is required to navigate through the model output current field in comparison to navigation in the absence of currents. The curve for straight line trajectories in the presence of currents, following colored squares in Fig. 6a, is shifted to the up and to the right with respect to the energy-optimal curve. Therefore, for every vehicle speed, the energy-optimal trajectory is saving both time and energy. Savings are increased for slower vehicles, and in this environment, time-savings increase more so than energy.

Fig. 6b illustrates four specific trajectories for $\|\mathbf{v}_0\| = [0.5, 1.0, 1.5, 2.0]$ m/s. In each case, the vehicle deviates from the straight-line path to follow the average circulation of the eddy. The fastest vehicles follow relatively direct paths, while the slowest vehicle which has a nominal speed similar to the current velocity traces the outer edge of the eddy where the currents are strongest. This trajectory is examined in more detail in Fig. 7 where it is compared to additional planning strategies. Time series of speed and power are shown in Figs. 7d-e, normalized by $\|\mathbf{v}_0\|$ and its associated power consumption, respectively. The direction of travel is referenced relative to the course, with $0°$ indicating the straight line from start to end, and positive angles increasing in the counterclockwise direction.

The vehicle following the energy-optimal trajectory moves slower than the time-optimal trajectory, as seen in Fig. 7d. This slower speed results in a less direct route, allowing the vehicle to exploit more of the structure in the current field, as shown in Fig. 7a. Figs. 7b-d highlight each vehicle's control input in response to the current field. As both optimal strategies allow control over the vehicle's direction, the optimal control aligns the vehicle's heading with the direction of the current, which alternates across the eddy. This alignment is especially pronounced in the energy-optimal strategy, where the vehicle slows down as it becomes more aligned with the current. These findings are consistent with the optimal control policy derived in Section III-A.

Finally, Table II summarizes the integrated results of time and energy consumption over the full trajectory, compared to the baseline line-following strategy. Both optimal strategies outperform line-following by more than 10% in either time or

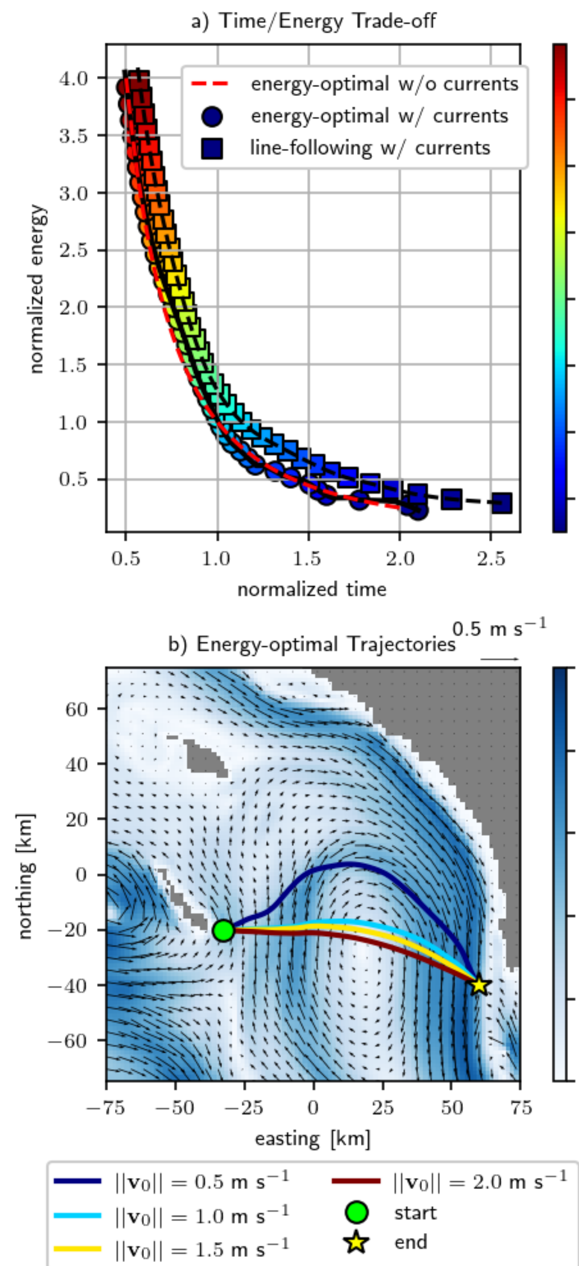

Fig. 6. a: Total time vs. energy for a range of vehicle parameterizations (color) following energy-optimal (circles) and line-following (squares) paths with a reference curve (red) for the solution without currents. Time and energy are normalized such that a vehicle with $\|\mathbf{v}_0\| = 1.0$ m/s completes the trajectory in unit time and energy. b: Select transects according to $\|\mathbf{v}_0\| = [0.5, 1.0, 1.5, 2.0]$ m/s are plotted on top of the time-averaged current field for the duration of the longest trajectory. Trajectories are colored by their no-current optimal speed.

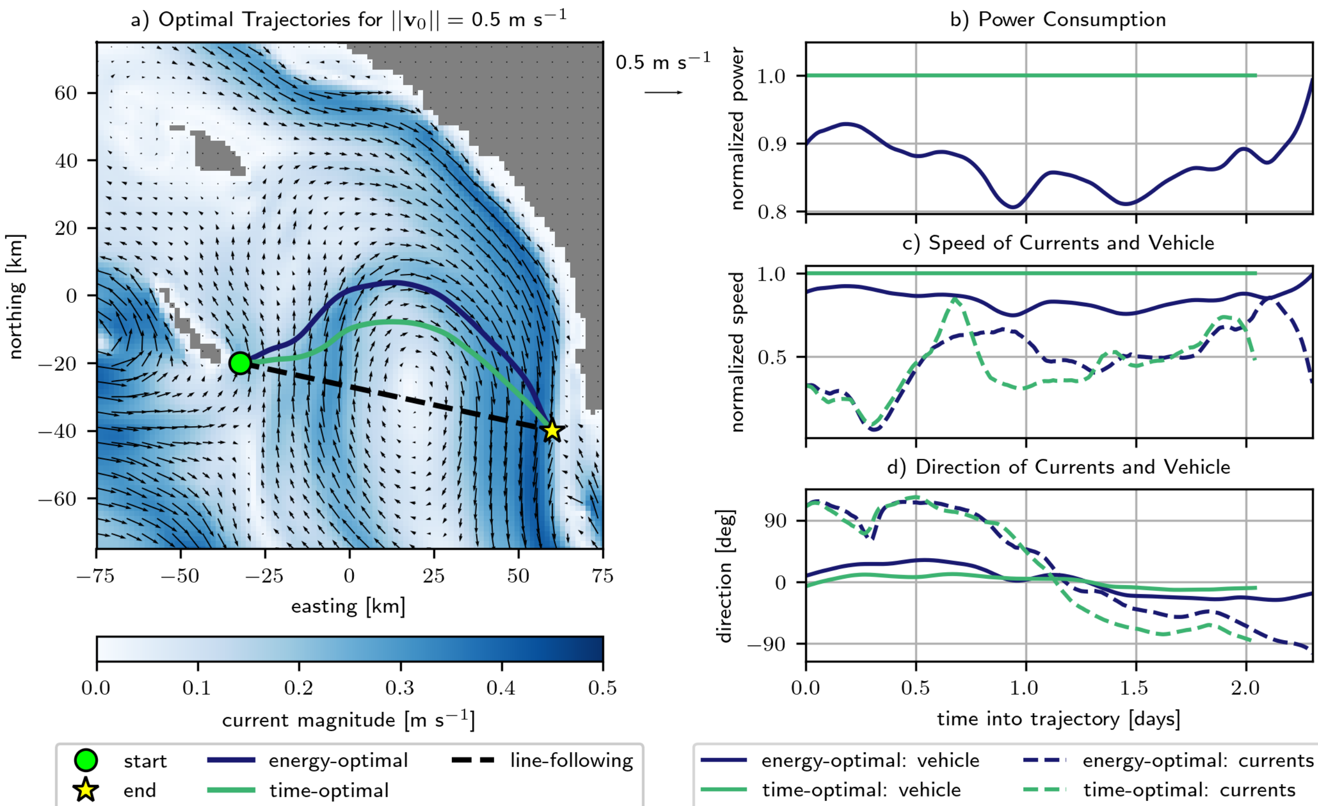

Fig. 7. a: Energy-optimal, time-optimal, and line-following trajectories for a vehicle with a no-current optimal speed $\|\mathbf{v}_0\| = 0.50$ m/s in the Southern California Bight are plotted on top of the time-averaged current field for the duration of the longest trajectory. b-d: Time series of power consumption, speed, and direction of each vehicle along its trajectory are shown by solid lines. Speed and direction of the currents interpolated along the respective trajectories and are denoted by dashed lines of the same color. Speed and power are normalized by their no-current optimal values, respectively. Direction is relative to the direction from $\mathbf{x}_0$ to $\mathbf{x}_f$.

TABLE II
SAVINGS RELATIVE TO LINE-FOLLOWING FOR THE TIME-OPTIMAL AND ENERGY-OPTIMAL TRAJECTORIES SHOWN IN FIG. 7.

Optimization	**Time Savings**	**Energy Savings**
Energy-optimal	10.1%	21.8%
Time-optimal	20.3%	20.3%

energy. The energy-optimal strategy proves to be the most efficient, though it only offers a 1.5% relative energy savings over the time-optimal strategy, while requiring 10% more time to complete the transit. In practice, an operator might opt for the time-optimal, fixed-speed solution depending on operational considerations.

B. *Ram Head, St. John*

In the second case study, we focus on a shorter transit near Ram Head, St. John where the time-dependent characteristics of the flow will have greater effects on the resulting optimal path. Operating autonomous vehicles near headlands or island wakes is difficult due to the strong, dynamic flows generated around abrupt topography. However, data collection in these areas is highly valuable due to increased biological activity [56] and their contribution to energy dissipation through mixing [57]. The local coordinate system is centered at 18° 17.99’ N, 64° 42.14’ W, and the starting position is $\mathbf{x}_0^* = (-1.75, -0.75)$ km and final position is $\mathbf{x}_f^* = (-0.50, 0.00)$ km. At this location, we consider no-current optimal speeds over the range $\|\mathbf{v}_0\| = [0.15, 1.25]$ m/s. To initialize Algorithm 1 we use a grid of initial conditions where $\mathcal{T}_f$ spans

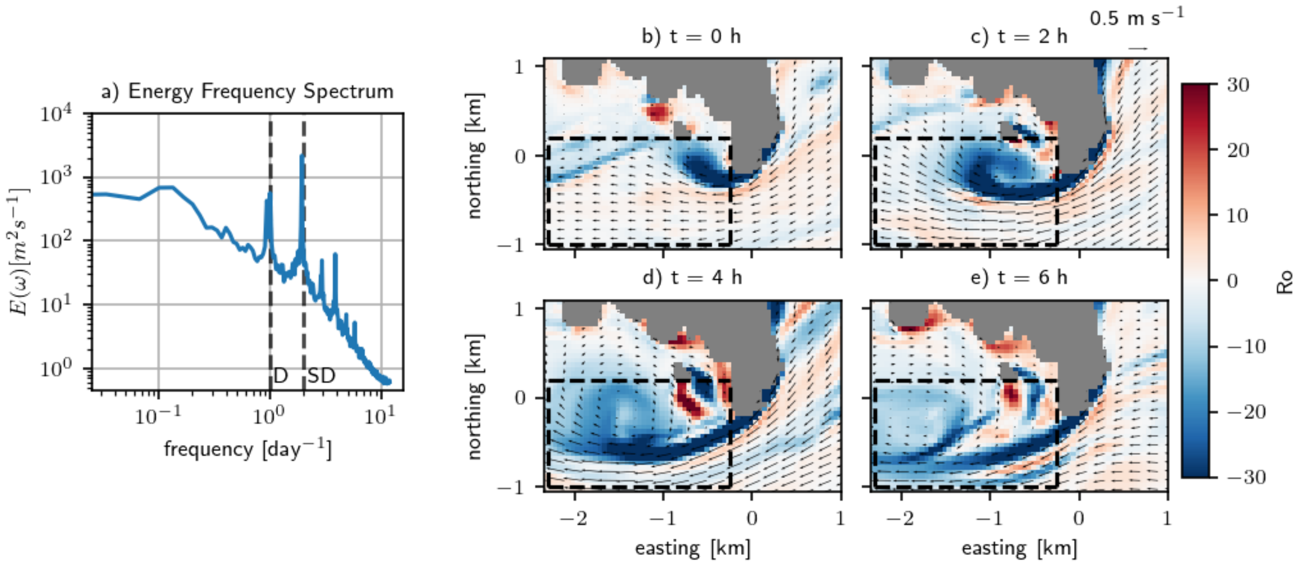

Fig. 8. a: Spectral characteristics of a tidal vortex generated at Ram Head, Virgin Islands, USVI. Kinetic energy frequency spectra are computed at each grid cell of the boxed region and averaged in space. Vertical lines annotate diurnal (D), and semi-diurnal (SD) frequencies, respectively. b-e: The temporal evolution of the vortex over 6 hours generated at the southern point and advecting westward. Maps colored by Rossby number (Ro) are overlaid with quiver arrows for surface currents. Positive and negative Ro represent cyclonic and anti-cyclonic circulation, respectively.

$[0\text{h}, 6\text{h}]$ with 0.1 h spacing and Θ_f spans $[-90°, 90°]$ with $2°$ spacing. The integrator used a 1-minute time step.

The ocean current field is obtained from a high-resolution Regional Ocean Modeling System (ROMS) run [58]. The model was run with hourly time steps, a horizontal resolution of 50 m, and 25 terrain following vertical levels, although, we focus only on the surface currents. This model output was previously used by McCammon et al. [59], which focused on developing methods for flow prediction from sparse observations. Here, we focus on the effects of a concentrated and dynamic feature on planning for slow-moving AUVs. Shown in Fig. 8, strong flow from east to west interacts with the southern tip of the island producing an anticyclonic vortex.

As in the section above, we first examine the paths of different vehicles, parameterized by $\|\mathbf{v}_0\|$. Each vehicle is initialized at the same time, corresponding to the flow shown in Fig. 8b. The time/energy tradeoff curve for the optimal trajectories follows the same general shape as the no-current curve, although there is a significant bias towards higher energy consumption, Fig. 9a. This is a result of a strong flow that accelerates around the point and is for the most part opposing the direction of travel, in the southern half of the domain. Northwest of the point there is a clockwise rotating flow which the slower vehicles use to advect along passively.

Selected trajectories, Fig. 9b, are shown for vehicles with no-current optimal speeds $\|\mathbf{v}_0\| = [0.15, 0.25, 0.5, 1.25]$ m/s. Once again we see that faster vehicles prefer more direct paths, while the slowest vehicles take a much longer and indirect path while using the currents of the advecting eddy. The optimal paths deform smoothly with the parameter $\|\mathbf{v}_0\|$. In comparison to the trajectories shown in Southern California, we highlight vehicles over a lower range of $\|\mathbf{v}_0\|$ to achieve large deviations from the direct path. The nominal transit here is much shorter than what was examined in Southern California, so slower vehicles benefit more from optimal planning. This highlights characteristic parameters of the problem; the relative scale of path length, vehicle speed, and current speed contribute to planning efficacy and complexity of the resulting trajectory.

Next, we examine the trajectories for $\|\mathbf{v}_0\| = 0.25$ m/s more closely, as shown in Fig. 10. In this environment, we only compare the energy-optimal and time-optimal trajectories because the line-following policy does not allow the vehicle to travel between the specified start and end points. Since the cross-

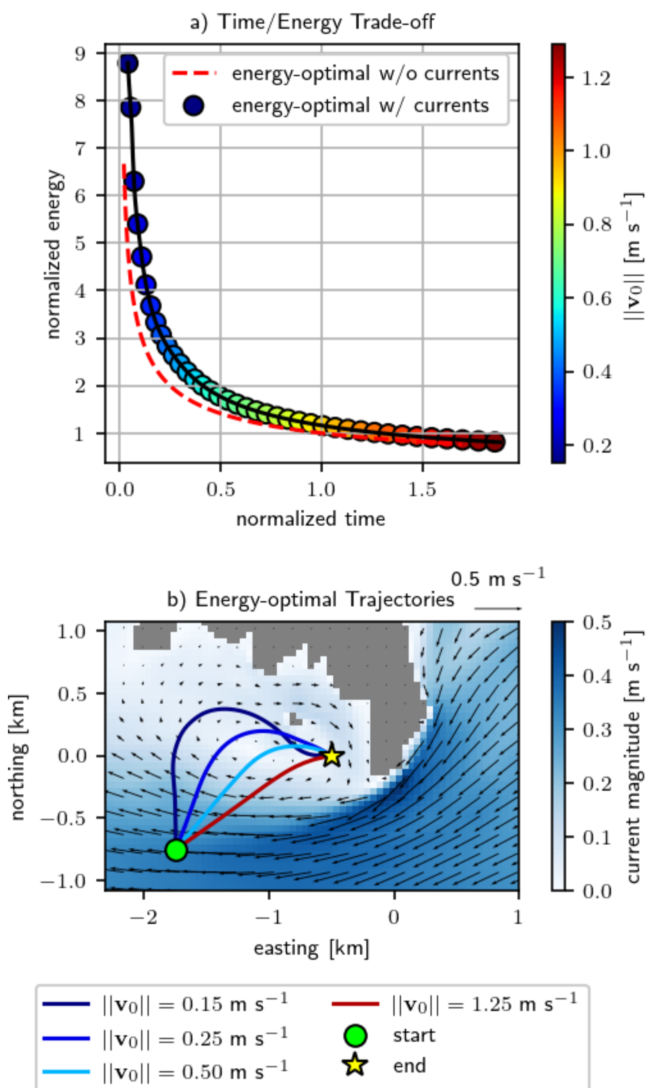

Fig. 9. a: Total time vs. energy for a range of vehicle parameterizations (colored circles) following energy-optimal paths with a reference curve (red) for the solution without currents. Time and energy are normalized such that a vehicle with $\|\mathbf{v}_0\| = 1.0$ m/s completes the trajectory in unit time and energy. b: Selec t trajectories corresponding to $\|\mathbf{v}_0\| = [0.15, .25, 0.5, 1.25]$ m/s are plotted on top of the time-averaged current field for the duration of the longest trajectory. Trajectories are colored by their no-current optimal speed.

track current is greater than 0.25 m/s the vehicle cannot maintain course with this control policy. In contrast to the previous example, the energy-optimal trajectory is shorter in distance and in time. In this current field, the ability to control speed dictates that the vehicle should go faster than the no-current optimal speed while it faces opposing currents, which is the case for the majority of its trajectory. Once the vehicle wraps around the northern side of the vortex, approximately 1.5 hours into the trajectory, the currents point in a favorable direction and it slows down, Fig. 10c. In the beginning of the trajectory, the fixed-speed vehicle orients its heading directly east canceling out the westward component of the current resulting in a course over ground directly northward. It only begins to move east once the vehicle is clearly in the lee of the island and uses the rotational component of the eddy to its advantage. In comparing the two control-strategies, the variable-speed trajectory is 22.0% faster and uses 6.8% less energy than the fixed-speed trajectory.

Due to the dynamic nature of this feature, we also examine the effects of a variable launch time t_0. As shown in Fig. 8a, tidal forcing dominates, with the majority of energy concentrated at the semidiurnal (SD) frequency. With the same geographic boundary conditions, we solve for energy-optimal trajectories for a vehicle with a no-current optimal speed of $\|\mathbf{v}_0\| = 0.15$ m/s launched within a 48 hr time window. By explicitly solving for the optimal trajectories at hourly launch times within this window, we aim to highlight the temporal correlation between physical processes in the environment and their effect on energy-consumption.

Alternatively, the maximum principle can be extended to directly optimize the start time. In this case, in addition to the canonical equations (11) and (12), the optimal control (17), and the transversality condition (18), optimal trajectories must also satisfy the following transversality condition at the free start time t_0

$$H(t_0) = 0. \tag{35}$$

The same boundary condition on the AUV position is enforced at a candidate launch time t_0. However, as we are interested in the time series of trajectories over multiple tidal cycles, we treat t_0 as fixed for each hour and solve for each trajectory using Algorithm 1.

The total energy consumption, total transit time, initial Hamiltonian $H(t_0)$, and free stream currents are shown in Fig. 11 for each launch time in the 48 hr window. The oscillations in the current's speed are reflected in the trajectory energy cost. The components of the free stream flow, Fig. 11d, are taken from a single grid point of the model located below the tip of the island at approximately $(0, -0.6)$ km. Over this time period the average magnitude of the free stream is 0.23 m/s with excursions exceeding

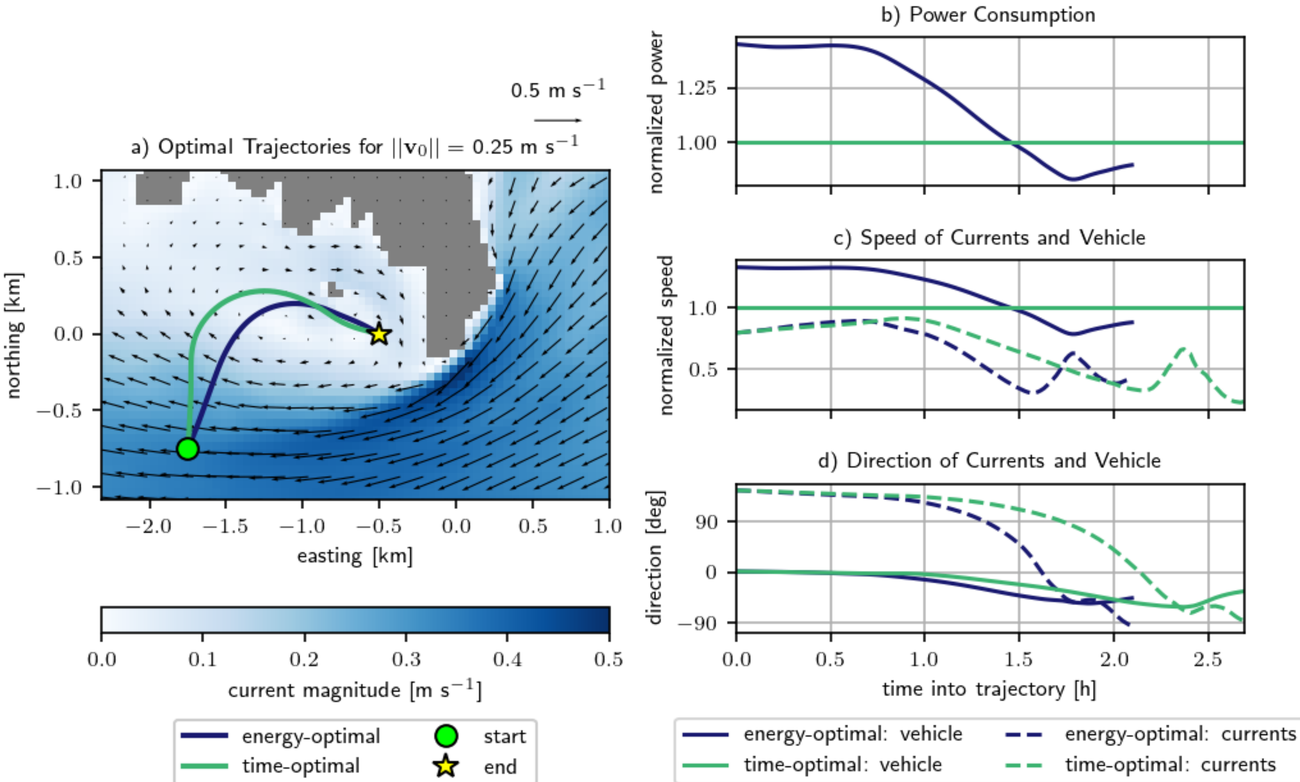

Fig. 10. a: Energy-optimal and time-optimal trajectories for a vehicle with a no-current optimal speed $\|\mathbf{v}_0\| = 0.25$ m/s near Ram Head, St. John are plotted on top of the time-averaged current field for the duration of the longest trajectory. b-d: Time series of power consumption, speed, and direction of each vehicle along its trajectory are shown by solid lines. Speed and direction of the currents interpolated along the respective trajectories and are denoted by dashed lines of the same color. Speed and power are normalized by their no-current optimal values, respectively. Direction is relative to the direction from $\mathbf{x}_0$ to $\mathbf{x}_f$.

0.45 m/s when the flow accelerates to the north west and subsequently causes vortex generation. Additionally, we see that the maximum values of E_f and t_f, respectively, occur in the hours prior to generation when the free stream flow is weakest and minimum values occur after generation when it is strongest. Notably, launching at the maximum energy start time consumes 95.5% more energy than at the minimum energy start time. Throughout the time series, the maxima and minima of energy consumption align with the zeros of $H(t_0)$ in agreement with (35).

Fig. 12 compares the two trajectories that result from the minimum and maximum energy start time, respectively. At the minimum energy start time, the vortex has already separated from the island tip and is advecting towards the starting location, while the vortex has just begun to form at the maximum energy start time. Due to this phase difference, the former can complete a majority of the trajectory while aligned with a component of the current where it is commanded to travel slower than $\|\mathbf{v}_0\|$ and consume much less power, as seen in Fig. 12c-d. On the other hand, the trajectory corresponding to the maximum energy start time travels much further into the lee of the island, while it waits for the vortex to move down stream, before it turns towards the end point. Over more than half of its trajectory, the AUV is traveling faster than $\|\mathbf{v}_0\|$. For both trajectories, there is an abrupt change in current direction, seen in Fig. 12e, which occurs at approximately 2.6 h and 4.5 h for minimum and maximum energy trajectories, respectively. This corresponds to the vehicle exiting the coherent circulation of the eddy. In each case, the vehicle is no longer traveling with the direction of strong currents, and the optimal control responds by increasing its velocity, seen in Fig. 12d.

VI. Summary and Future Work

In this work we applied PMP to derive the governing equations for energy-optimal trajectories

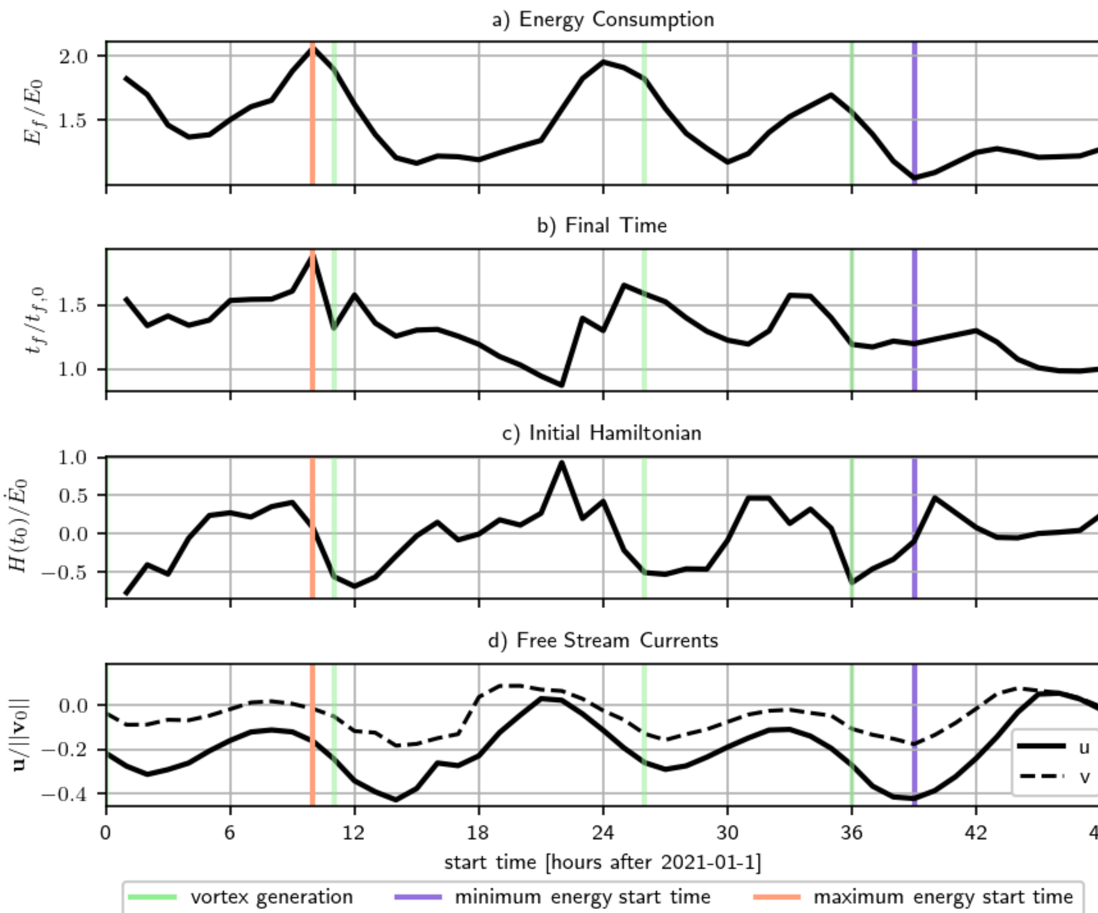

Fig. 11. a-d: Time series of energy consumption (a), travel time (b), Hamiltonian at t_0 (c), and components of free stream currents (d) as a function of transit start time. The no-current optimal speed of the vehicle in each simulation is $\|\mathbf{v}_0\| = 0.15$ m/s. Energy, time, and the Hamiltonian at t_0 are normalized by the time, energy, and power, respectively, of the optimal trajectory in a no-current environment. Vertical lines annotate the start time of the vehicle corresponding to minimum (purple) and maximum (orange) energy consumption, and gray lines correspond to times at the start of repeated vortex generation.

by an AUV in a continuous environment. This study highlighted the applicability of PMP for energy-efficient transit; however, due to its generality, our framework can be extended to multi-objective optimization that can consider risk or information gain, in addition to time and energy.

We used PMP to formulate a Hamiltonian system of ODEs for the vehicle and adjoint dynamics, along with boundary conditions that, when satisfied, provide a first-order condition for optimality. Through an analysis of the transversality condition, we gained insight into the optimal control by isolating the effects of the vehicle parameters and the currents, respectively. We found that a vehicle can be parameterized by its speed in the absence of currents, $\|\mathbf{v}_0\|$, which corresponds to various combinations of predetermined vehicle coefficients and hotel loads. Subsequently, we found that the scaling of the optimal speed, $\|\mathbf{v}^*\|/\|\mathbf{v}_0\|$, is only a function of the along-track component of the currents. We numerically solved the associated TPBVP in a straightforward manner with pure shooting and demonstrated its effectiveness in various flow fields. Our main insight was to use an analytical solution to the transversality condition to simplify the inputs to the shooting function. Instead of finding an initial condition on the adjoint we showed that the adjoint vector is determined by the final heading of an optimal trajectory. Therefore, we defined a shooting function that integrates backwards in time, away from the goal point and reduced the dimension of the search space. We illustrated our methods in three canonical current fields including a 3-D, time varying flow, and conducted case studies

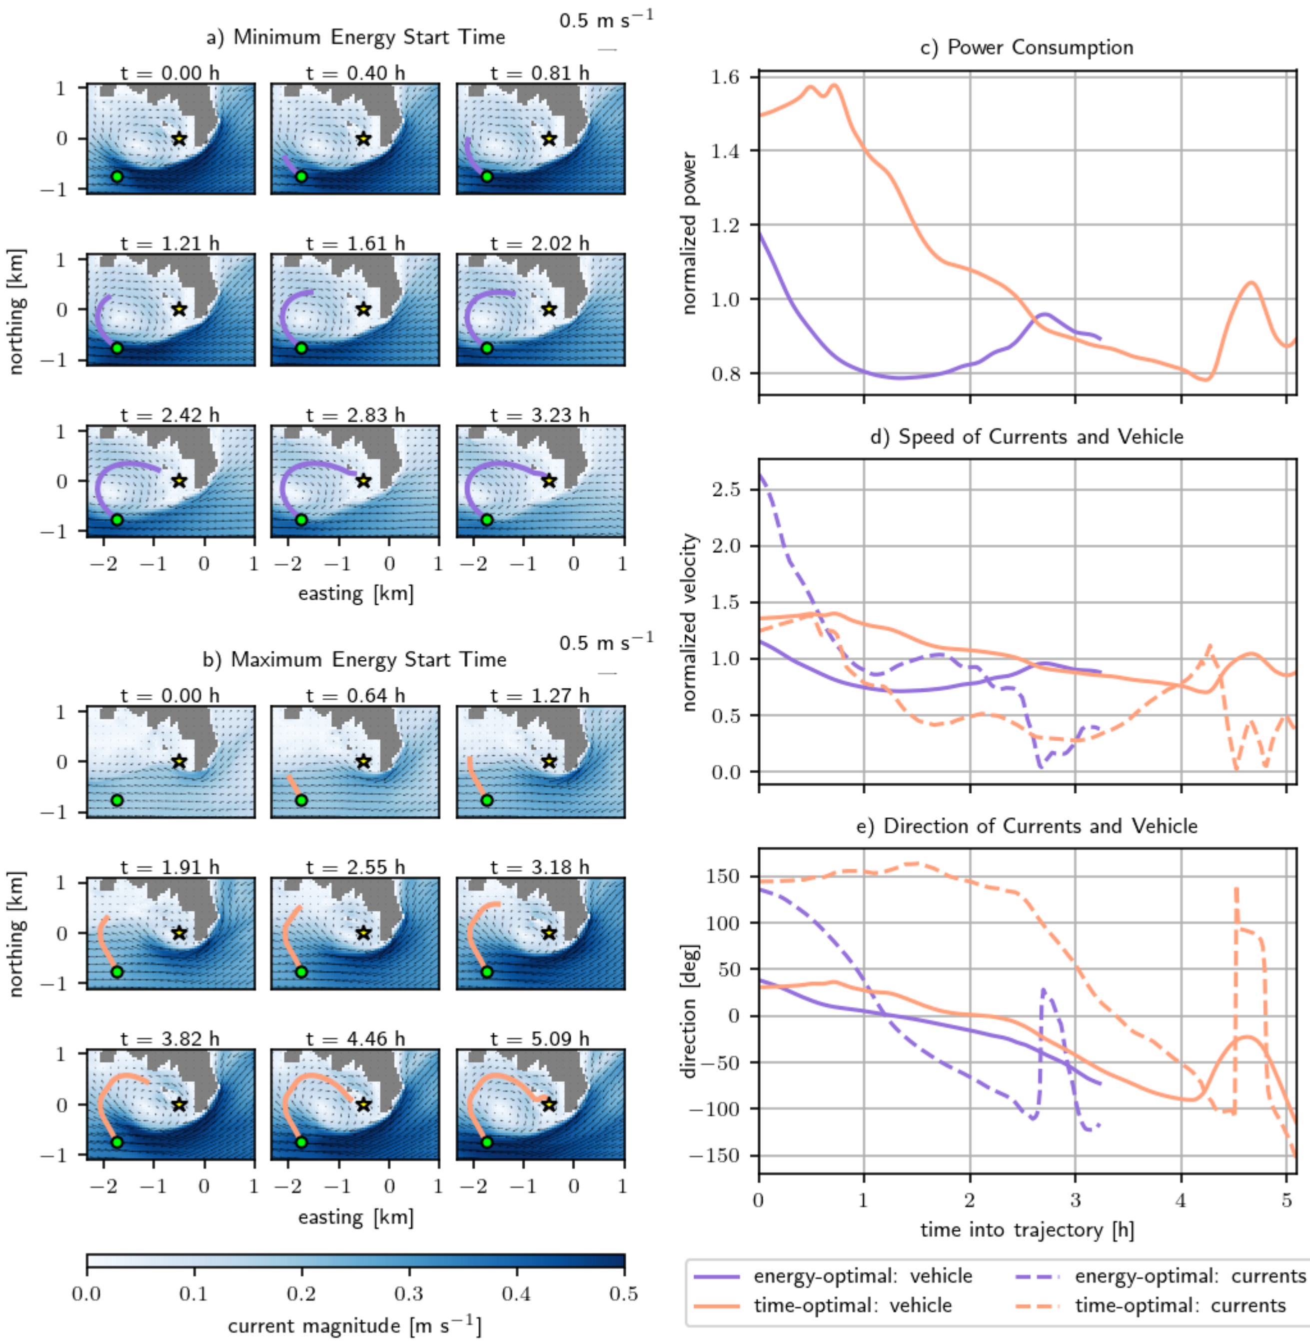

Fig. 12. a-b: Temporal evolution of the currents and energy-optimal trajectories at start times associated with minimum (a) and maximum (b) energy consumption from Fig. 11. In each panel, the start and end points are shown as green circles and yellow stars, respectively. c-e: Time series of power consumption, speed, and direction of each vehicle along its trajectory are shown by solid lines. Speed and direction of the currents interpolated along the respective trajectories and are denoted by dashed lines of the same color. Speed and power are normalized by their no-current optimal values, respectively. Direction is relative to the direction from $\mathbf{x}_0$ to $\mathbf{x}_f$.

in realistic currents from high-resolution, tidally-resolving ocean models.

In the Southern California case study, focused on long-duration transits where the currents are dominated by persistent features, we compared the performance of energy-optimal (variable-speed) and time-optimal (fixed-speed) trajectories. This analysis showed that the fixed-speed trajectory was over 10% faster for a marginal (1.5%) decrease in energy efficiency. In contrast, the Ram Head case study, which focused on a much more dynamic, small-scale feature, highlighted the importance of adaptive speed choices. In this case the energy-optimal trajectory provided over 20% time-savings along with 6% energy-savings. Additionally, we showed the importance of mission start-time by analyzing optimal trajectories over a 2-day window, where within 6 hours the total energy consumption fluctuates by over 50%. This range of energy consumption was attributed to semi-diurnal tidal forcing, which produced oscillatory ocean currents at the same frequency and generated the observed headland eddy.

PMP with pure shooting offers a simple alternative to well-established approaches that involve solving complex PDEs or discretizing the state space. Its appeal lies in its relative ease of implementation and potential to scale favorably with problem dimension. These computational benefits are particularly attractive for applications that require evaluating many optimal trajectories, such as finding an optimal launch time. However, PDE-based methods *guarantee* an optimal solution, so long as it exists. In certain environments, the chaotic nature of optimal trajectories derived from ODE-based methods [24] increase the difficulty of selecting a suitable initial condition, particularly in higher dimensions. High sensitivity to initial conditions can hinder convergence to the boundary conditions, potentially preventing a valid solution. As stated in Section III-C, future work can consider more robust numerical methods, such as multiple shooting, to directly replace the current root-finding procedure. It should also consider addressing this issue by developing hybrid methods. For example, building on the framework in [11], a coarse-grid PDE solution could capture the macro-scale influence of the current field on the trajectory, providing an informed initial condition for our method to refine. Machine-learning-based methods can also be explored to provide informed initial values for the adjoint and final final time [60], [61]. Alternatively, PDE-based approaches can be directly improved with adaptive mesh refinement [62].

Our methodology applies to the pre-planning stage of an autonomous deployment, and a key challenge when adapting to real-world situations is the reliance on forecast data. Our formulation assumes a deterministic environment where flow fields are perfectly known, but in practice, forecast accuracy degrades over time. One way to mitigate this issue is to re-plan when an updated forecast becomes available. However, this approach depends on frequent communication with a shore-based planner and requires the AUV to surface. Another approach could be to reformulate the problem in terms of stochastic optimal control, and minimize the *expected* energy consumption for which alternative methods exist [63]. Given the ability to produce flow fields that are statistically representative of a given environment or forecast, our methodology is well suited to evaluate the distribution of many trajectories and their energy cost. However, neither of these approaches explicitly incorporate in situ data from the vehicle, which could validate or invalidate the forecast model. If observations confirm the forecast, following the planned trajectory remains beneficial, but if discrepancies arise, reverting to a more conservative planning strategy may be preferable. For a more flexible solution, future work can address onboard data assimilation to update a forecast online, allowing real-time re-optimization of the planned trajectory. We hope that our method can be integrated with onboard algorithms to build a more complete and robust autonomous system.

VII. Acknowledgments

The authors thank Matthew Mazloff of Scripps Institution of Oceanography for providing MITgcm model output of the Southern California Bight. The authors also thank Seth McCammon and Gordon Zhang of Woods Hole Oceanographic Institution for providing ROMS model output for Ram Head, St. John. The authors would like to thank the Office of Naval Research for their enduring support of undersea research. Raymond Young was supported by the National Defense Science and Engineering Graduate fellowship.

APPENDIX A
STATEMENT OF PMP FOR A FIXED-ENDPOINT, FREE-TIME OPTIMAL CONTROL PROBLEM

This appendix provides the full statement of PMP which is used to derive the methods in Section III-A. The original proof was developed by Pontryagin et al. [7], but we follow the notation of Liberzon [10]. We require the statement of PMP used to solve a time independent fixed-endpoint, free-time OCP. The addition of time dependence can be addressed by appending t to the state vector. This class of OCP is defined for a control system defined by a system of ODEs

$$\dot{x} = f(x, u), \qquad x(t_0) = x_0, \tag{36}$$

where $x = x(t) \in \mathbb{R}^n$, the state, and $u = u(t) \in \mathcal{U}$, the control, are both functions of time. $\mathcal{U}$ is the control set and t_0 is the initial time. The goal of the OCP is to minimize a cost functional $J(u)$,

$$J(u) = \int_{t_0}^{t_f} L(t, x(t), u(t))dt, \tag{37}$$

where $L : \mathbb{R} \times \mathbb{R}^n \times \mathcal{U} \mapsto \mathbb{R}$ is the running cost, and t_f, the final time, is defined indirectly through a target set. The target set S is defined as

$$S \subset [t_0, \infty) \times \{x_f\}, \tag{38}$$

where x_f is the prescribed goal state and t_f is the smallest time such that $(t_f, x(t_f)) \in S$. Additionally we require that f, $\partial f/\partial x$, L, and $\partial L/\partial x$ are continuous.

The maximum principle states that if there exists an optimal control $u^* : [t_0, t_f] \to \mathcal{U}$ and corresponding state trajectory $x^* : [t_0, t_f] \to \mathbb{R}^n$, then there is a function $p^* : [t_0, t_f] \to \mathbb{R}^n$ and a constant $p_0^* \leq 0$ satisfying $(p_0^*, p^*(t)) \neq (0, 0)$ for all $t \in [t_0, t_f]$ and having the following properties:

1) Canonical equations: the dynamics of x^* and p^* are expressed as the Jacobian of a Hamiltonian H

$$\dot{x}^* = \frac{\partial H}{\partial p}(x^*, u^*, p^*, p_0^*), \tag{39}$$

$$\dot{p}^* = -\frac{\partial H}{\partial x}(x^*, u^*, p^*, p_0^*) \tag{40}$$

with boundary conditions $x^*(t_0) = x_0$ and $x^*(t_f) = x_f$. The Hamiltonian $H : \mathbb{R}^n \times U \times \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ is a function of the state, control, and adjoint, and is defined as

$$H(x, u, p, p_0) := \langle p, f(x, u)\rangle + p_0 L(x, u) \tag{41}$$

2) Optimal control via maximization of the Hamiltonian: Along an optimal trajectory the function $u \mapsto H(x^*(t), u, p^*(t), p_0^*)$ has a global maximum at $u = u^*(t)$. In other words, the control maximizes the Hamiltonian and satisfies the inequality

$$\begin{aligned} &H(x^*(t), u^*(t), p^*(t), p_0^*) \\ &\qquad \geq H(x^*(t), u, p^*(t), p_0^*) \end{aligned} \tag{42}$$

which holds for all $t \in [t_0, t_f]$ and all $u \in U$.

3) Transversality condition:

$$H(x^*(t_f), u^*(t_f), p^*(t_f), p_0^*) = 0. \tag{43}$$

In the time independent case,

$$H(x^*(t), u^*(t), p^*(t), p_0^*) = 0. \tag{44}$$

for all $t \in [t_0, t_f]$.

In our case, (p^*, p_0^*) can be normalized so that $p_0^* = 1$.

APPENDIX B
APPLICATIONS OF PMP TO RELATED PROBLEMS

A. Time-optimal trajectories

This methodology is used for all examples in section IV and in both case studies from Section V. Time-optimal trajectories are equivalent to energy-optimal trajectories when $\|\mathbf{v}(t)\| = v_{\max}$ for all t. Therefore, the dynamics (1) can be simplified to:

$$\dot{\mathbf{x}}(t) = v_{\max}\hat{\mathbf{v}}(t) + \mathbf{u}(t, \mathbf{x}), \tag{45}$$

where the updated control set is $\mathcal{V} = \mathbb{S}^d$. The OCP can also be simplified with the running cost $L(t) = 1$:

$$\hat{\mathbf{v}}^*(t) = \underset{\hat{\mathbf{v}}(t) \in \mathcal{V}, t_f \in \mathbb{R}^+}{\arg\min} \int_{t_0}^{t_f} 1\, dt. \tag{46}$$

Now we can follow the procedure in Section III-A and invoke the three statements of PMP:

1) Definition of the Hamiltonian and canonical equations:

$$H(t, \mathbf{x}, \mathbf{p}) = v_{\max}\langle \hat{\mathbf{v}}, \mathbf{p}\rangle + \langle \mathbf{u}, \mathbf{p}\rangle - 1, \tag{47}$$

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}} = v_{\max}\hat{\mathbf{v}} + \mathbf{u}, \tag{48}$$

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}} = -\nabla_{\mathbf{x}}\mathbf{u}^{\mathrm{T}}\mathbf{p}. \tag{49}$$

We remark that the evolution of the adjoint is equivalent in both the time-optimal and energy-optimal problem.

2) Maximization of the Hamiltonian: Now there is no control over the magnitude of $\mathbf{v}$, and the same argument from the energy-optimal problem applies to maximize H based on $\hat{\mathbf{v}}$:

$$\hat{\mathbf{v}}^*(t) = \hat{\mathbf{p}}(t). \tag{50}$$

3) Transversality Condition: The transversality condition $H(t, \mathbf{x}(t_f), \mathbf{p}(t_f) = 0$ remains the same. In order to initialize the adjoint vector we assume an optimal final direction $\hat{\mathbf{v}}(t_f) = \hat{\mathbf{p}}(t_f)$ and determine its magnitude:

$$v_{\max}\langle \hat{\mathbf{p}}, \mathbf{p}\rangle + \langle \mathbf{u}, \mathbf{p}\rangle - 1 = 0 \tag{51}$$

$$\|\mathbf{p}\| = \frac{1}{v_{\max} + u_\theta} \tag{52}$$

The solution procedure and numerical implementation remains the same.

B. Energy-optimal trajectories w/ time constraint

One simple extension to the energy-optimal OCP presented in the main text is to enforce a constraint on the arrival time t_f. This extension is motivated by Doshi et al. [20] in which the goal is to find the Pareto-optimal family of solutions with respect to time and energy.

In our formulation, this is equivalent to changing the target set to $S = \{t_f\} \times \{\mathbf{x}_f^*\}$. As a consequence, the transversality condition no longer holds, and $H(t_f, \mathbf{x}(t_f), \mathbf{p}(t_f) = C$ for an unknown constant C [10]. There is one fewer boundary condition, however, t_f is now and fixed. Therefore, there is enough information for a well defined shooting function:

$$S_3 : \begin{pmatrix} H(t_f) \\ \hat{\mathbf{v}}(t_f) \end{pmatrix} \mapsto \left(\mathbf{x}(t_0) - \mathbf{x}_0^*\right). \tag{53}$$

Now, the magnitude of the adjoint, at t_f, can be initialized as the solution to

$$\frac{2}{3}\sqrt{\frac{1}{3\alpha}}\|\mathbf{p}(t_f)\|^{3/2} + u_\theta(t_f, \mathbf{x}_f^*, \hat{\mathbf{v}}^*(t_f))\|\mathbf{p}(t_f)\| - \beta = H(t_f). \tag{54}$$

The same numerical procedure can be followed when the search over the grid $\mathcal{T}_f$ is replaced by a grid of potential values of $H(t_f)$. Furthermore, if we define $\beta' = \beta - H(t_f)$ we see that these solutions are equivalent to augmenting the artificially augmenting the hotel load, which causes the vehicle to speed up.

REFERENCES

[1] Z. Zeng, L. Lian, K. Sammut, F. He, Y. Tang, and A. Lammas, "A survey on path planning for persistent autonomy of autonomous underwater vehicles," *Ocean Engineering*, vol. 110, pp. 303–313, Dec. 2015.

[2] J. Hwang, N. Bose, and S. Fan, "AUV Adaptive Sampling Methods: A Review," *Applied Sciences*, vol. 9, no. 15, 2019.

[3] C. Cheng, Q. Sha, B. He, and G. Li, "Path planning and obstacle avoidance for AUV: A review," *Ocean Engineering*, vol. 235, p. 109355, Sept. 2021.

[4] L. Lu, H. Wang, N. Gu, S. Li, Z. Peng, and D. Wang, "Research on Path Planning of Underwater Vehicle with Optimal Energy Consumption," in *2024 39th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, pp. 2222–2226, June 2024.

[5] Q. Nguyen, L. Huang, and A. Perera, "Energy optimal UAV path planning for full area coverage: A Review," in *2025 9th International Conference on Robotics, Control and Automation (ICRCA)*, pp. 218–222, Mar. 2025.

[6] R. Bellman, "The theory of dynamic programming," *Bulletin of the American Mathematical Society*, vol. 60, no. 6, pp. 503–515, 1954.

[7] L. S. Pontryagin, *Mathematical Theory of Optimal Processes*. London: Routledge, Oct. 2017.

[8] F. H. Clarke and R. B. Vinter, "The Relationship between the Maximum Principle and Dynamic Programming," *SIAM Journal on Control and Optimization*, vol. 25, pp. 1291–1311, Sept. 1987.

[9] P. Cannarsa and H. Frankowska, "Some Characterizations of Optimal Trajectories in Control Theory," *SIAM Journal on Control and Optimization*, vol. 29, pp. 1322–1347, Nov. 1991.

[10] D. Liberzon, *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, 2012.

[11] E. Cristiani and P. Martinon, "Initialization of the Shooting Method via the Hamilton-Jacobi-Bellman Approach," *Journal of Optimization Theory and Applications*, vol. 146, no. 2, pp. 321–346, 2010.

[12] T. Lolla, M. P. Ueckermann, K. Yigit, P. J. Haley, and P. F. J. Lermusiaux, "Path planning in time dependent flow fields using level set methods," in *2012 IEEE International Conference on Robotics and Automation*, pp. 166–173, IEEE, May 2012.

[13] T. Lolla, P. F. J. Lermusiaux, M. P. Ueckermann, and P. J. Haley, "Time-optimal path planning in dynamic flows using level set equations: theory and schemes," *Ocean Dynamics*, vol. 64, pp. 1373–1397, Oct. 2014.

[14] T. Lolla, P. J. Haley, and P. F. J. Lermusiaux, "Time-optimal path planning in dynamic flows using level set equations: realistic applications," *Ocean Dynamics*, vol. 64, pp. 1399–1417, Oct. 2014.

[15] T. Lolla, P. J. Haley Jr., and P. F. J. Lermusiaux, "Path planning in multi-scale ocean flows: Coordination and dynamic obstacles," *Ocean Modelling*, vol. 94, pp. 46–66, Oct. 2015.

[16] D. N. Subramani, P. F. J. Lermusiaux, P. J. Haley, C. Mirabito, S. Jana, C. S. Kulkarni, A. Girard, D. Wickman, J. Edwards, and J. Smith, "Time-optimal path planning: Real-time sea exercises," in *OCEANS 2017 - Aberdeen*, pp. 1–10, June 2017.

[17] C. S. Kulkarni and P. F. J. Lermusiaux, "Three-dimensional time-optimal path planning in the ocean," *Ocean Modelling*, vol. 152, p. 101644, Aug. 2020.

[18] D. N. Subramani and P. F. J. Lermusiaux, "Energy-optimal path planning by stochastic dynamically orthogonal level-set optimization," *Ocean Modelling*, vol. 100, pp. 57–77, Apr. 2016.

[19] D. N. Subramani, P. J. Haley Jr., and P. F. J. Lermusiaux, "Energy-optimal path planning in the coastal ocean," *Journal of Geophysical Research: Oceans*, vol. 122, no. 5, pp. 3981–4003, 2017.

[20] M. M. Doshi, M. S. Bhabra, and P. F. Lermusiaux, "Energy–time optimal path planning in dynamic flows: Theory and schemes," *Computer Methods in Applied Mechanics and Engineering*, vol. 405, p. 115865, Feb. 2023.

[21] J. Brandman and C. Olson, "Globally Time-Optimal Path Planning for Unmanned Underwater Vehicles in Three-Dimensional Current Fields Using Hamilton-Jacobi Partial Differential Equations," in *OCEANS 2023 - MTS/IEEE U.S. Gulf Coast*, (Biloxi, MS, USA), pp. 1–10, IEEE, Sept. 2023.

[22] C. Y. Kao, S. Osher, and J. Qian, "Lax–Friedrichs sweeping scheme for static Hamilton–Jacobi equations," *Journal of Computational Physics*, vol. 196, pp. 367–391, May 2004.

[23] C.-W. Shu, "High order numerical methods for time dependent hamilton-jacobi equations," in *Mathematics and Computation in Imaging Science and Information Processing*, vol. 11 of *Lecture Notes Series, Institute for Mathematical Sciences, National University of Singapore*, pp. 47–91, WORLD SCIENTIFIC, Oct. 2007.

[24] L. Biferale, F. Bonaccorso, M. Buzzicotti, P. Clark Di Leoni, and K. Gustavsson, "Zermelo's problem: Optimal point-to-point navigation in 2D turbulent flows using reinforcement learning," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, p. 103138, Oct. 2019.

[25] R. E. Davis, N. E. Leonard, and D. M. Fratantoni, "Routing strategies for underwater gliders," *Deep Sea Research Part II: Topical Studies in Oceanography*, vol. 56, pp. 173–187, Feb. 2009.

[26] B. Liebchen and H. Löwen, "Optimal navigation strategies for active particles," *Europhysics Letters*, vol. 127, p. 34003, Sept. 2019.

[27] V. N. Dobrokhodov, C. Walton, I. I. Kaminer, and K. D. Jones, "Energy-Optimal Guidance of Hybrid Ultra-Long Endurance UAV," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 15639–15646, 2020.

[28] F. L. Pereira, T. Grilo, and S. Gama, "Optimal Control Framework for AUV's Motion Planning in Planar Vortices Vector Field," in *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, pp. 1–6, Nov. 2018.

[29] B. Garau, A. Alvarez, and G. Oliver, "Path Planning of Autonomous Underwater Vehicles in Current Fields with Complex Spatial Variability: an A* Approach," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pp. 194–198, Apr. 2005.

[30] D. Rao and S. B. Williams, "Large-scale path planning for Underwater Gliders in ocean currents," in *Australasian Conference on Robotics and Automation (ACRA)*, (Sydney, Australia), 2009.

[31] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, Dec. 1959.

[32] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, pp. 100–107, July 1968. Conference Name: IEEE Transactions on Systems Science and Cybernetics.

[33] B. Garau, M. Bonet, A. Alvarez, S. Ruiz, and A. Pascual, "Path Planning for Autonomous Underwater Vehicles in Realistic Oceanic Current Fields: Application to Gliders in the Western Mediterranean Sea," *Journal of Maritime Research*, vol. 6, no. 2, pp. 5–22, 2009.

[34] T.-B. Koay and M. Chitre, "Energy-efficient path planning for fully propelled AUVs in congested coastal waters," in *2013 MTS/IEEE OCEANS - Bergen*, pp. 1–9, June 2013.

[35] D. Kularatne, S. Bhattacharya, and M. Ani Hsieh, "Time and Energy Optimal Path Planning in General Flows," in *Robotics: Science and Systems XII*, Robotics: Science and Systems Foundation, 2016.

[36] D. Kularatne, S. Bhattacharya, and M. Hsieh, "Going with the flow: a graph based approach to optimal path planning in general flows," *Autonomous Robots*, vol. 42, Oct. 2018.

[37] J. Chang, M. Anderson, S. Merrifield, A. Nager, R. Hess, R. Young, S. Kitchen, and E. Terrill, "Power Efficiency Autonomy for Long Duration AUV Operation," in *2022 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*, pp. 1–6, Sept. 2022.

[38] M. Anderson, A. Nager, S. Merrifield, E. Terrill, R. Young, and J. Chang, "Analysis and Test of an Optimal Propeller Speed Schedule to Reduce AUV Energy Consumption," in *OCEANS 2025 Brest*, pp. 1–8, June 2025.

[39] J. G. Bellingham, Y. Zhang, J. E. Kerwin, J. Erikson, B. Hobson, B. Kieft, M. Godin, R. McEwen, T. Hoover, J. Paul, A. Hamilton, J. Franklin, and A. Banka, "Efficient propulsion for the Tethys long-range autonomous underwater vehicle," in *2010 IEEE/OES Autonomous Underwater Vehicles*, pp. 1–7, Sept. 2010.

[40] J. Sherman, R. Davis, W. Owens, and J. Valdes, "The autonomous underwater glider "Spray"," *IEEE Journal of Oceanic Engineering*, vol. 26, pp. 437–446, Oct. 2001.

[41] R. E. Davis, C. Eriksen, and C. Jones, "Autonomous Buoyancy-Driven Underwater Gliders," *Technology and Applications of Autonomous Underwater Vehicles*, pp. 37–58, Nov. 2002.

[42] J. Manley and S. Willcox, "The Wave Glider: A persistent platform for ocean science," in *OCEANS'10 IEEE SYDNEY*, pp. 1–5, May 2010.

[43] A. Amador, S. T. Merrifield, R. A. McCarthy, R. Young, and E. J. Terrill, "Wave Glider Speed Model for Real-time Motion Planning," in *OCEANS 2021: San Diego – Porto*, pp. 1–9, Sept. 2021.

[44] C. L. Gentemann, J. P. Scott, P. L. F. Mazzini, C. Pianca, S. Akella, P. J. Minnett, P. Cornillon, B. Fox-Kemper, I. Cetinić, T. M. Chin, J. Gomez-Valdes, J. Vazquez-Cuervo, V. Tsontos, L. Yu, R. Jenkins, S. D. Halleux, D. Peacock, and N. Cohen, "Saildrone: Adaptively sampling the marine environment," *Bulletin of the American Meteorological Society*, vol. 101, no. 6, pp. E744–E762, 2020.

[45] M. M. Doshi, M. S. Bhabra, and P. F. J. Lermusiaux, "Energy–time optimal path planning in dynamic flows: Theory and schemes," *Computer Methods in Applied Mechanics and Engineering*, vol. 405, p. 115865, Feb. 2023.

[46] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 3 ed., 2007.

[47] J. Stoer and R. Bulirsch, "Ordinary Differential Equations," in *Introduction to Numerical Analysis* (J. Stoer and R. Bulirsch, eds.), pp. 404–535, New York, NY: Springer, 1980.

[48] U. M. Ascher, R. M. M. Mattheij, and R. D. Russell, *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*. Society for Industrial and Applied Mathematics, 1995.

[49] J. Kierzenka and L. F. Shampine, "A BVP solver based on

residual control and the Maltab PSE," *ACM Transactions on Mathematical Software*, vol. 27, pp. 299–316, Sept. 2001.

[50] S. K. Lam, A. Pitrou, and S. Seibert, "Numba: a LLVM-based Python JIT compiler," in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, LLVM '15, (New York, NY, USA), pp. 1–6, Association for Computing Machinery, Nov. 2015.

[51] R. Young, S. Merrifield, M. Anderson, and E. Terrill, "A Simulation Framework for Environmentally-Aware Autonomous Underwater Vehicle (AUV) Mission Planning and Algorithm Development," in *OCEANS 2023 - MTS/IEEE U.S. Gulf Coast*, pp. 1–6, Sept. 2023.

[52] D. L. Rudnick, K. D. Zaba, R. E. Todd, and R. E. Davis, "A climatology of the California Current System from a network of underwater gliders," *Progress in Oceanography*, vol. 154, pp. 64–106, May 2017.

[53] C. Hill and J. Marshall, "- Application of a parallel Navier-Stokes model to ocean circulation," in *Parallel Computational Fluid Dynamics 1995* (A. Ecer, J. Periaux, N. Satdfuka, and S. Taylor, eds.), pp. 545–552, Amsterdam: North-Holland, Jan. 1996.

[54] M. R. Mazloff, B. Cornuelle, S. T. Gille, and J. Wang, "The Importance of Remote Forcing for Regional Modeling of Internal Waves," *Journal of Geophysical Research: Oceans*, vol. 125, Feb. 2020.

[55] R. Young, S. Merrifield, M. Anderson, M. Mazloff, and E. Terrill, "A Greedy Depth-Seeking Behavior for Energy-Efficient Transits by an Autonomous Underwater Vehicle," *IEEE Journal of Oceanic Engineering*, vol. 49, pp. 1383–1396, Oct. 2024.

[56] W. G. Zhang, Y. Jia, A. Apprill, and T. A. Mooney, "Fine-scale hydrodynamics around St. John, U.S. Virgin Islands. Part I: spatial and temporal heterogeneity in the coastal environment," *Frontiers in Marine Science*, vol. 12, June 2025.

[57] T. S. Johnston, M. C. Schönau, T. Paluszkiewicz, J. A. MacKinnon, B. K. Arbic, P. L. Colin, M. H. Alford, M. Andres, L. Centurioni, H. C. Graber, K. R. Helfrich, V. Hormann, P. F. Lermusiaux, R. C. Musgrave, B. S. Powell, B. Qiu, D. L. Rudnick, H. L. Simmons, L. St. Laurent, E. J. Terrill, D. S. Trossman, G. Voet, H. W. Wijesekera, and K. L. Zeiden, "Flow Encountering Abrupt Topography (fleat): A Multiscale Observational and Modeling Program to Understand How Topography Affects Flows in the Western North Pacific," *Oceanography*, vol. 32, no. 4, pp. 10–21, 2019.

[58] A. F. Shchepetkin and J. C. McWilliams, "The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean Modelling*, vol. 9, pp. 347–404, Jan. 2005.

[59] S. McCammon, Y. Jia, and W. G. Zhang, "Predicting Time-Varying Flow Fields Using Sparse Observations with Self-Organizing Maps," in *2024 IEEE OES AUV Symposium*, 2024.

[60] Y. Zang, J. Long, X. Zhang, W. Hu, W. E, and J. Han, "A machine learning enhanced algorithm for the optimal landing problem," in *Proceedings of Mathematical and Scientific Machine Learning* (B. Dong, Q. Li, L. Wang, and Z.-Q. J. Xu, eds.), vol. 190 of *Proceedings of Machine Learning Research*, pp. 319–334, PMLR, 15–17 Aug 2022.

[61] J. Shi, J. Wang, L. Su, Z. Ma, H. Chen, J. Shi, J. Wang, L. Su, Z. Ma, and H. Chen, "A Neural Network Warm-Started Indirect Trajectory Optimization Method," *Aerospace*, vol. 9, Aug. 2022.

[62] Z. Li and P. Song, "Adaptive mesh refinement techniques for the immersed interface method applied to flow problems," *Computers & structures*, vol. 122, pp. 249–258, June 2013.

[63] D. N. Subramani, Q. J. Wei, and P. F. J. Lermusiaux, "Stochastic time-optimal path-planning in uncertain, strong, and dynamic flows," *Computer Methods in Applied Mechanics and Engineering*, vol. 333, pp. 218–237, May 2018.

Raymond Young received the B.S. and M.S. degrees in mechanical engineering from the University of California San Diego, in 2020 and 2022, respectively. He received the Ph.D. degree in oceanography from Scripps Institution of Oceanography, University of California, San Diego, in 2025.

He is currently a postdoctoral scholar with Scripps Institution of Oceanography where his research interests include autonomous marine vehicles, optimal control, and robotic path planning.

F. Hunter Akins received a B.A. in applied mathematics and a B.A. in physics from the University of California, Berkeley in 2017. He received the Ph.D. degree in oceanography from Scripps Institution of Oceanography, University of California, San Diego, in 2023. His research is on underwater acoustic signal and array processing and their application to experimental acoustic data.

Sophia T. Merrifield received the Ph.D. degree in physical oceanography from the Massachusetts Institute of Technology/Woods Hole Oceanographic Institution Joint Program, in 2016.

She is currently an acting Associate Professor with Scripps Institution of Oceanography, University of California San Diego. Her research interests include marine vehicle autonomy and oceanic turbulence.