

Coverage-Maximizing Solar-Powered Autonomous Surface Vehicle Control for Persistent Gulf Stream Observation

Kavin Govindarajan, Ben Haydon, Kirti Mishra, and Chris Vermillion¹

Abstract—The Gulf Stream, which comes within 100 km of the United States coastline in both the Florida Straits and the vicinity of Cape Hatteras, is estimated to possess over 160 TWh/year of technical energy capacity. To better understand the behavior of the Gulf Stream, whose flow resource varies in both space and time, a relatively sparse network of fixed acoustic Doppler current profilers (ADCPs) and shore-mounted high-frequency radar units have been supplemented by more granular but infrequent boat transect runs and undersea glider deployments. Collectively, these measurements provide highly granular data with respect to either time or space, but not both. This work represents part of a comprehensive effort to evaluate use of a solar-powered autonomous surface vehicle (ASV) fleet to supplement existing observational capabilities. The proposed solar-powered ASV can provide data with high spatial and temporal granularity, but comes with the challenge of optimally planning its mission in an adaptive manner. To address this challenge in this work, we propose a multilevel controller that fuses the A* search algorithm with an upper level waypoint selector and lower level heading control. Focusing on a critically important mission domain adjacent to Cape Hatteras, and relying on a Mid-Atlantic Bight, South Atlantic Bight Regional Ocean Model (MAB-SAB-ROM), we compare the performance of our proposed algorithm against several competing strategies. We demonstrate a significant performance improvement in terms of a dynamic coverage metric, both in comparison to competing strategies and to the existing observational network.

I. INTRODUCTION

The Gulf Stream possesses over 160 TWh/year of technical energy capacity in the region between Florida and North Carolina [1]. This has led to millions of dollars in research aimed at characterizing the spatial and temporal variations in the Gulf Stream profile, including the North Carolina Renewable Ocean Energy Program, PEACH Project, Southeast National Marine Renewable Energy Center, and other efforts.

To facilitate the process of characterizing spatiotemporal variations in the Gulf Stream profile adjacent to Cape Hatteras, several observational tools have been employed. These include moored acoustic Doppler current profilers (ADCPs), boat-mounted ADCPs, a high-frequency radar network, and undersea gliders. Specifically, the numerous (10+ in the vicinity of Cape Hatteras, although not all concurrently) moored ADCPs provide measurements of current velocity vs. depth at 4-meter intervals, averaged over each hour [2]. These provide excellent temporal and depth-wise spatial resolution; however, they provide little capacity for spatial

extrapolation to latitudes and longitudes where the stationary ADCPs are not deployed. Boat transect runs, where research vessels are equipped with their own ADCPs [3], provide a spatially granular picture of the current profile over the approximately 24-hour period of the transect run. However, cost and personnel restrictions result in practical limitations on how often these runs can be made; in practice, they are typically made no more frequently than annually, leading to very limited temporal resolution. To provide a more granular spatiotemporal surface characterization of the Gulf Stream, a high-frequency radar network has been deployed [4], providing spatial data every 6 km and temporal data every hour. However, these data are limited to surface measurements and are not nearly as accurate as on-site ADCP measurements. Finally, several undersea gliders have been deployed for the purpose of providing longer-duration spatiotemporal characterizations of the Gulf Stream at depth [5], [6]. However, their limited velocities (usually less than 1 km/hour of velocity made good) limit the gliders to downstream paths, and generally no more than approximately one month in duration.

Despite the extensiveness of the Gulf Stream observational network, coverage remains spatiotemporally sparse. Temporally granular measurements like moored ADCPs are restricted to a fixed spatial location, and spatially granular measurements like boat transects provide limited temporal resolution. Persistent Autonomous Surface Vehicles (ASVs) with downward-looking ADCPs, such as sailing drones [7], wave gliders [8], and solar-powered ASVs [9], can help to close the aforementioned gaps by providing continuous operations over months at a time and a large spatial domain. In this work, we will consider the SeaTrac SP-48 ASV shown in Fig. 1. This system addresses the major concerns with the variety of other solutions:

- The ASV is a mobile system allowing for measurements across the entire domain.
- As an autonomous system, the ASV is significantly cheaper to operate than a full research vessel.
- Unlike an undersea glider, the ASV's solar panels enable *persistent missions*, i.e. missions whose length is not capped by physical constraints.
- The ASV's approximately 4.5 knot peak velocity allows it to make upstream progress outside of the center ("jet") of the Gulf Stream.

Although autonomous mobile systems address the main shortcomings of other approaches, they introduce a control challenge of their own: optimal real-time exploration of

¹Kavin Govindarajan, Ben Haydon, Kirti Mishra, and Chris Vermillion serve in the roles of Undergraduate Research Assistant, Graduate Research Assistant, Postdoctoral Research Associate, and Associate Professor, respectively, with the Department of Mechanical and Aerospace Engineering, North Carolina State University, Raleigh, NC 27606, USA.




Fig. 1: SeaTrac ASV considered in this work [9] – Image used with permission.

TABLE I: SeaTrac ASV Parameters

Variable Name	Value	Units
Length	4.8	meters
Beam	1.39	meters
Draft	0.42	meters
Weight	300	kilograms
Solar Panel Area	4.17	meters ²
Brushless Motor	500	Watts
Battery Capacity	6.75	kWh
Top Speed	4.5	kts
Cruise Speed	3	kts

the mission domain in light of uncertain mobility arising from the solar resource and spatiotemporally varying current. The standard mechanism for controlling such systems is a so-called *line transect* strategy, wherein the ASV executes pre-planned “lawn mower”-type patterns in order to cover the domain. While such strategies indeed traverse the full mission domain, they do so without consideration of the propulsive resource and other disturbances, leading typically to sub-optimal paths. Adaptive sampling literature has examined more sophisticated information-driven planning approaches for general oceanographic and atmospheric observation research, which have included dynamic coverage-based strategies [10], [11], [12], exploration/exploitation strategies [13], [14], [15], [16], [17], and other information-driven approaches [18], [19], [20], [21] to account for the spatiotemporally varying propulsion resource and disturbances in performing real-time planning.

It is the goal of the present research to apply information-driven planning to the control of the SeaTrac SP-48 ASV over a critical Gulf Stream mission domain. We propose a hierarchical approach where an upper-level mission planner computes a coverage-optimal target waypoint within the global mission domain, and a lower-level planner computes the path to that waypoint. Specifically, two implementations of the lower-level planner are investigated here: (i) a direct-to-point strategy and (ii) an A* algorithm [22]. We choose the A* algorithm due to its use of heuristic parameters, enabling us to tune the algorithm to prioritize the desired heuristics. We compare the results of the proposed optimal hierarchical strategy to a line transect approach, which is a commonly used standard approach for observational studies.

II. MODELING

A. Dynamic Models

In this work, we consider the SeaTrac SP-48, which is an autonomous, propeller-driven autonomous surface vessel (ASV) with an on-board battery and solar panels for recharging. Given that the time scales of the environment greatly exceed those of the boat dynamics, we assume direct control of the boat’s velocity and heading according to the following dynamic model:

$$\begin{aligned}\dot{x} &= v_b \cos \theta_b + v_f(x, y, t) \cos \theta_f(x, y, t) \\ \dot{y} &= v_b \sin \theta_b + v_f(x, y, t) \sin \theta_f(x, y, t)\end{aligned}\quad (1)$$

where x and y are the two spatial states of the system, v_b and θ_b are the boat speed and boat heading (which serve as control variables), respectively, and $v_f(x, y, t)$ and $\theta_f(x, y, t)$ are the spatiotemporally varying flow speed and direction at the x, y location of the boat. We also model the state of charge of the on-board battery, denoted by b , according to:

$$\begin{aligned}\dot{b} &= \eta_s A_s I(t) - \eta_m F_D \|\vec{v}_b\|_2 - P_e \\ F_D &= \frac{1}{2} \rho A_w C_D \|\vec{v}\|_2^2\end{aligned}\quad (2)$$

where η_s is the solar panel efficiency, A_s is the solar panel area, $I(t)$ is the solar shortwave radiation (in units of W/m^2), η_m is the motor efficiency, F_D is the drag force on the boat, and P_e is the electrical power consumption of the on-board electronic systems. The drag force, F_D , is modeled as a function of the wetted area, A_w , the drag coefficient, C_D , and the *apparent* current speed, \vec{v} . In order to simulate a system with these dynamics, it is necessary to characterize the environmental variables $I(t)$ and $\vec{v}_f(x, y, t)$, in addition to defining a control architecture for determining \vec{v}_b (which characterizes both the heading and speed of the boat). The former requires the selection of a meaningful mission domain and the acquisition of appropriate solar and current data over that domain. The latter requires the definition of a control architecture. We discuss the approaches to each task in the following sections.

B. Site Selection

For our mission domain, we selected the region shown in Fig. 2 off the coast of North Carolina frequently occupied by the Gulf Stream. This site was selected for two main reasons:

- 1) This section of the Gulf Stream has been shown to be nearest to the shore, making it a promising location for the future installation of marine hydrokinetic energy-harvesting devices [2].
- 2) This section of the Gulf Stream has also been estimated to possess the smallest level of spatial variability due to so-called *meanders*. Nevertheless, existing observational data is insufficient to precisely pin down the exact resource potential and variability in this area.

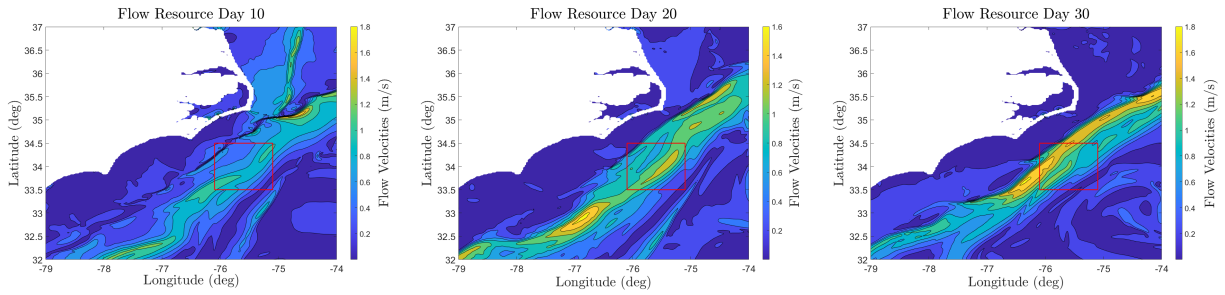


Fig. 2: Mission domain and flow resource (at the surface, based on the MAB-SAB-ROM model) at 10 day intervals

C. Flow Data

To characterize the surface currents for the purpose of driving our dynamic simulations, we have elected to use the Mid-Atlantic Bight, South Atlantic Bight Regional Ocean Model (MAB-SAB-ROM) [23]. This model outputs the flow speed vector as a function of the 3-dimensional x, y, z location and time (t). Because the SP-48 is an ASV, we only rely on the surface elements of the data (i.e., those corresponding to $z = 0$) to drive the dynamics of the ASV.

D. Solar Data

Solar data from the ERA-Interim model [24], available as a function of x, y location and time (t), was used to drive the ASV dynamics. Because the length scales of the solar data are significantly longer than the dimensions of the mission domain, we assume in this work that the solar resource is temporally but not spatially varying.

E. Dynamic Coverage

The scientific objective of ASV-based Gulf Stream observations lies in the persistent coverage of a mission domain whose characteristics (e.g., current velocity as a function of depth and x, y location) are known to change over time. Thus, for the development and subsequent validation of our proposed control strategies, it is essential that we rely on a metric that captures the quality with which the algorithms achieve persistent coverage over the domain. To formulate this metric, we turn to the concept of *dynamic coverage*, discussed in [10], [12], and [25], which characterizes the quality with which a mobile agent characterizes a domain. This characterization typically depends on the definition of a *sensing function*, which characterizes the quality of information available at each point within the domain as a function of the mobile agent(s)'s location(s) – points far away from or unobservable by an agent will be associated with low sensing function values.

In our work, we rely on a mathematical coverage model with two key features that are representative of the problem at hand:

- In the absence of available measurements near a particular x, y location, the coverage at that location gradually declines.
- Coverage at a particular x, y location immediately increases when the ASV passes sufficiently close to that

location, where the level of increase depends on the proximity of the ASV to that x, y location.

These key characteristics are captured in our work through the following discrete-time coverage model:

$$q(x, y, t + \delta t) = \max \{q(x, y, t) - q_{loss}\delta t, S(d(t))\}, \quad (3)$$

where q is dynamic coverage, q_{loss} is a flat coverage loss rate, $d(t)$ is the distance at time t between the ASV and the point (x, y) about which coverage is being calculated, and $S(d(t))$ is the *sensing function*, given by:

$$S(d(t)) = e^{-\frac{d(t)^2}{2l^2}}. \quad (4)$$

Here, l is a sensing length scale, which can roughly be thought of as a radius within which the ASV is able to accumulate coverage. Examining equation (3), one can see that if the ASV is far from the (x, y) point under consideration, coverage at that location will continue to decrease. However, when the ASV gets sufficiently close to an (x, y) point, coverage will once again increase in accordance with the sensing function. Note that the left term in equation (3) is always decreasing and the right term is bounded above by 1, so coverage at a point will also be bounded above by 1.

In this work, we are interested in maximizing a *spatially averaged* measure of coverage over the entire mission domain. To characterize this, we focus on the following average coverage measure in both the control formulations and subsequent analysis of the results:

$$J(t) = \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m q(x(i), y(j), t), \quad (5)$$

where n is the number of discrete points in the x dimension, and m is the number in the y dimension. The quantity $J(t)$ represents the average coverage at time t , which is bounded above by 1 and below by 0, with higher values representing more coverage, i.e., more knowledge of the environment.

III. CONTROL FRAMEWORK

In this work, we propose a multi-level control strategy, which can be depicted according to the block diagram in Fig. 3. We compare two versions of the strategy: one in which only an upper-level coverage-based waypoint selector and path-following controller is implemented, and another where the upper-level waypoint is fed into a coverage-based A* controller that generates waypoints for the path-following




Fig. 3: Block diagram of the control approach.

controller. Furthermore, each of these algorithms is compared against a baseline line transect strategy (which executes prescribed “lawn mower” patterns within the mission domain). Descriptions of the individual blocks within the hierarchical control structure are given in the following subsections.

A. Upper-Level Waypoint Selection

The upper level waypoint selector selects waypoints for the vehicle according to the following optimization:

$$[x_{goal}, y_{goal}] = \arg \max_{(x_i, y_j) \in D} \left(\frac{\Delta q(x_i, y_j)}{\Delta t(x_i, y_j)} \right) \quad (6)$$

where x_i and y_j are the decision variables belonging to the mission domain D ,

$$\Delta t(x_i, y_j) = \frac{\sqrt{(x_i - x_0)^2 + (y_j - y_0)^2}}{v_b}$$

represents the estimated time for the boat to travel from its current position (x_0, y_0) to (x_i, y_j) , and $\Delta q(x_i, y_j)$ is an estimate of the coverage that will be gained in traversing a straight line path from (x_0, y_0) to (x_i, y_j) that is defined as follows:

$$\Delta q(x_i, y_j) = \sum_{k=1}^N (1 - q_k - L_k) \quad (7)$$

where k is an index of each discrete point along the line between the current position (x_0, y_0) and (x_i, y_j) , N is the total number of points on that line, q_k is the current coverage at the point with index k , and

$$L_k = d \frac{q_{loss}}{v_b} \left(1 - \frac{k}{N} \right),$$

where $q_{loss} = 0.01$ represents the coverage loss rate, $d = \sqrt{(x_i - x_0)^2 + (y_j - y_0)^2}$ is the distance to the next potential waypoint, and L_k represents the amount of coverage that will be lost between the time the ASV reaches point k and the time it reaches the next potential waypoint (x_i, y_j) .

B. Local Coverage Maximization

The local coverage maximization controller utilizes the A* search algorithm to minimize the cost associated with traversal to the next waypoint. The A* cost associated with

traveling from point $p_1 = (x_1, y_1, t)$ to $p_2 = (x_2, y_2, t + \delta t)$ is defined as:

$$C(p_1, p_2) = k_1 (\Delta d_{goal}(p_1, p_2)) + k_2 (q(p_2) + F) \quad (8)$$

where $\Delta d_{goal}(p_1, p_2)$ represents the progress towards the final waypoint, which represents the goal location for A*. It is calculated as:

$$\Delta d_{goal}(p_1, p_2) = d_{goal}(p_1) - d_{goal}(p_2) \quad (9)$$

$$d_{goal}(p_i) = \sqrt{(x_{goal} - x_i)^2 + (y_{goal} - y_i)^2}. \quad (10)$$

Returning back to equation (8), $q(p_2)$ represents the coverage at point p_2 and

$$F = v_f \cdot \frac{1 - \cos(|\theta_f - \theta_{true}|)}{2}$$

represents the “usefulness” of the flow, where a flow direction aligned with the direction towards the goal results in a lower cost, and flow directions opposing the direction towards the goal resulting in a higher cost. This is augmented by the velocity of the flow, with high flow speeds in the direction towards the goal being most desirable, and thus resulting in a lower cost. The cost function is weighted according to k_1 and k_2 , where $k_1 = 0.05$ represents the weight on the term incentivizing progress towards final waypoint and $k_2 = 0.95$ represents a weight on the combined coverage and flow cost metric.

C. Path-Following Control

The path-following control block in each of the candidate control strategies is identical. The purpose of the block is to take a waypoint from an upper-level waypoint selector and speed from the upper-level speed controller as inputs and control the boat heading in order to reach the next waypoint. This controller selects the heading such that the true velocity points directly at the next waypoint. In equation form:

$$\theta_b = \theta_{true} + \sin^{-1} \left(\frac{v_f}{v_b} \sin(\theta_f - \theta_{true}) \right) \quad (11)$$

where θ_b is the calculated boat heading, θ_{true} is the heading towards the selected waypoint, θ_f is the direction in which the flow acts, and v_f, v_b are the flow velocity and boat velocity respectively. Note that this method would require modification if the boat velocity was slower than the maximum flow velocity since there would be no heading angle that would result in net progress towards a waypoint directly upstream.

D. Speed Control

Each of the candidate control systems uses the same global speed controller. The speed selected is in order to achieve zero net charging or discharging of the battery over a period corresponding to the traversal from one upper-level waypoint to the next. This is calculated as:

$$v_b = \|\vec{v}\|_2 = \sqrt[3]{\frac{(\eta_s A_s I_{avg} - P_e) \eta_{mp}}{\frac{1}{2} \rho A_w C_D}} \quad (12)$$

where each of the parameters is listed in Table II. This equation was derived from the battery dynamics (2). Future work will examine the incorporation of speed optimization strategies that strategically charge or discharge the battery according to the spatiotemporal forecast of the current and temporal characteristics of the solar resource. Properly implemented, this level of additional optimization is expected to further improve coverage.

TABLE II: Speed Control Parameters

Variable Name	Symbol	Value	Units
Solar Panel Area	A_s	4.17	m^2
Average Solar Shortwave Radiation	I_{avg}	175	W/m^2
On-board Electronics	P_e	100	W
Wetted Area	A_w	5.82	m^2
Solar Panel Efficiency	η_s	0.18	-
Motor + Propeller Efficiency	η_{mp}	0.25	-
Drag Coefficient	C_D	0.0030	-

IV. RESULTS AND DISCUSSION

In this section, we compare the results of the hierarchical coverage-based planning strategy of Section IV against two comparison strategies, both of which can be represented as simplified special cases of the block diagram of Fig. 3.

A. Comparison Strategy: Direct-to-Point Coverage-Based Control

The direct-to-point coverage-based strategy represents a simplified version of the control strategy of Fig. 3 for comparison purposes. Specifically, this variant of the controller retains the global waypoint selection but does not include the A*-based refinement. Instead, this strategy travels directly between waypoints selected by the upper-level waypoint selector. Thus, in this strategy, $x_i = x_{goal}$ and $y_i = y_{goal}$.

B. Comparison Strategy: Line Transect Strategy

The transect strategy represents our baseline control strategy. In the context of Fig. 3, the global waypoint selection and A*-based refinement are replaced with pre-selected waypoints (corresponding to x_i and y_i). These waypoints are selected in order to traverse the mission domain in an orderly manner, without regard for the spatiotemporally varying current or temporally varying solar resource (much like a lawnmower - see Fig. 8).

C. Simulation Results and Discussion

Spatially averaged coverage results from simulations of each strategy operating a 2-month long mission are displayed in 4. The simulation was conducted with the boat and environment model updated on 1-minute time-steps. The coverage-based hierarchical and direct-to-point strategies both outperformed the traditional line transect strategy due to their ability to directly account for flow conditions. Because each simulation starts with a coverage value of zero at all points within the mission domain, the first several days of each simulation are spent converging to a relatively steady level of coverage. Beyond this point, fluctuations in the

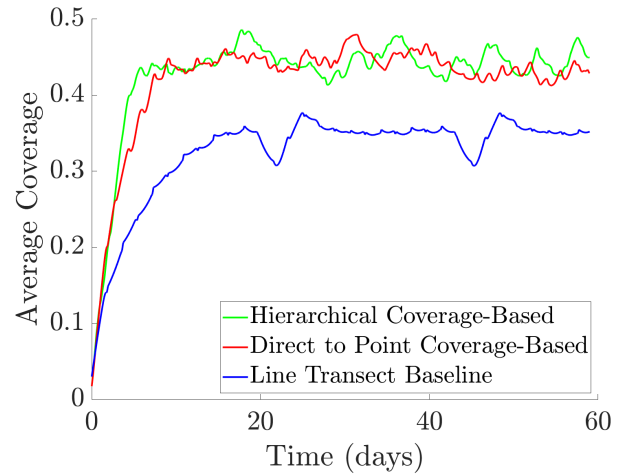


Fig. 4: Comparison of spatially averaged coverage vs. time under the three candidate control approaches.




Fig. 5: Comparison of spatially averaged coverage vs. time under the three candidate control approaches with slower boat velocity.

available solar resource and ocean current, along with periods in which the ASV traverses a recently visited point, result in persistent fluctuations. We examine the performance of the three algorithms in terms of both the transient period and (more importantly for persistent missions) the post-transient period.

It can first be seen from Fig. 4 that the rate at which coverage initially increases during the transient period is higher with more sophisticated algorithms. Specifically, the direct-to-point coverage-based strategy outperforms the line transect baseline, and the hierarchical coverage-based strategy outperforms both comparison strategies in terms of convergence speed. After the transient period, both coverage-based strategies continue to outperform the line transect baseline. While the hierarchical coverage-based strategy continues to show overall superior performance when averaged over the post-transient period, it is noteworthy that as cov-




Fig. 6: Path taken by Hierarchical Coverage-Based Controller alongside map of coverage in mission domain at days 15 and 60.




Fig. 7: Path taken by Direct to Point Coverage-Based Controller alongside map of coverage in mission domain at days 15 and 60.




Fig. 8: Path taken by Line Transect Baseline alongside map of coverage in mission domain at days 15 and 60.

erage approaches 45%, the direct-to-point strategy is more competitive with the hierarchical strategy during this phase. It is hypothesized that the ability of the simpler direct-to-point strategy to achieve similar performance over much of the post-transient period suggests that the “regularity” of the straight-line paths can lead to performance improvements in the long term. Considering this long-term behavior is a topic of ongoing research, either through model predictive control or an added term to the A* algorithm corresponding to a measure of path regularity.

To gain further insight into the relative performances of the candidate control strategies, it is instructive to consider (i) the actual paths traversed under each strategy, (ii) the available solar and flow resources as a function of time (noting that the 2D current velocity profiles at selected

“snapshots” in time are given in Fig. 2), and (iii) 2D coverage profiles at selected “snapshots” in time under each strategy. The paths traversed by the ASV are shown in Figs. 6, 7, and 8. As expected, the direct-to-point strategy engages in sweeping straight-line motions between highly disparate points in the domain, whereas the hierarchical strategy engages in excursions when those excursions are deemed beneficial for coverage. Temporal variation of the solar resource, in conjunction with spatiotemporal variation in the surface current velocity (available in discrete “snapshots” in Fig. 2) helps to explain some of the dips in coverage experienced during the post-transient period. Because of the spatial variability in the current, along with the fact that coverage improvement can only be attained when traversing a “new” area (or one not recently traversed), the timing and

extent of average coverage fluctuations will also depend on the location of the ASV. Finally, Figs. 6, 7, and 8 show a coverage map over the full 2D mission domain under all three candidate control approaches, at two “snapshots” in time. It can be seen here, at both snapshots in time, that in addition to maintaining higher average coverage, the hierarchical strategy does the best job of reducing the level of coverage variability across the domain. Furthermore, areas of extremely minimal coverage are small at each snapshot in time. In contrast, the line transect strategy results in large swaths of the mission domain having near-zero coverage at both time instances.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we performed a data-driven characterization of a hierarchical, coverage-based algorithm for persistent Gulf Stream exploration using a solar-powered ASV. The proposed algorithm was compared against a simplified direct-to-point coverage-based approach and a standard line transect approach. Simulation results, which were based on a MAB-SAB-ROM model for the ocean current and an ERA-Interim solar resource model, show that the hierarchical coverage-based algorithm outperforms the comparison algorithms in terms of (i) the attained average coverage level, (ii) the convergence time to a “steady” coverage level, and (iii) the consistency of coverage over the mission domain. Furthermore, the direct-to-point coverage-based algorithm was found to significantly outperform the line transect algorithm and approach the performance of the hierarchical algorithm over portions of the simulation. Future work will focus on augmenting the proposed coverage-based approach with an intelligent velocity planner that strategically charges and discharges the battery in order to attain faster or slower velocities than those achievable through the net-zero charge/discharge approach employed in this work.

REFERENCES

- [1] K. Haas, “Assessment of energy production potential from ocean currents along the United States coastline,” Georgia Inst. of Technology, Atlanta, GA (United States), Tech. Rep., 2013.
- [2] R. He, J. Bane, M. Muglia, S. Haines, C. Lowcher, Y. Gong, and P. Taylor, “Gulf Stream marine hydrokinetic energy resource characterization off Cape Hatteras, North Carolina USA,” in *OCEANS 2016 - Shanghai*, 2016, pp. 1–4.
- [3] M. Muglia, H. Seim, and P. Taylor, “Gulf Stream marine hydrokinetic energy off Cape Hatteras, North Carolina,” *Marine Technology Society Journal*, vol. 54, no. 6, pp. 24–36, 2020.
- [4] S. Haines, H. Seim, and M. Muglia, “Implementing quality control of high-frequency radar estimates and application to Gulf Stream surface currents,” *Journal of Atmospheric and Oceanic Technology*, vol. 34, no. 6, pp. 1207 – 1224, 2017. [Online]. Available: <https://journals.ametsoc.org/view/journals/atot/34/6/jtech-d-16-0203.1.xml>
- [5] G. Gawarkiewicz, R. E. Todd, W. Zhang, J. Partida, A. Gangopadhyay, M.-U.-H. Monim, P. Fratantoni, A. M. Mercer, and M. Dent, “The changing nature of shelf-break exchange revealed by the OOI Pioneer array,” *Oceanography*, vol. 31, no. 1, pp. 60–70, 2018.
- [6] R. E. Todd, “Export of Middle Atlantic Bight shelf waters near Cape Hatteras from two years of underwater glider observations,” *Journal of Geophysical Research: Oceans*, vol. 125, no. 4, p. e2019JC016006, 2020.
- [7] 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. Tsonetos, 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. [Online]. Available: <https://journals.ametsoc.org/view/journals/bams/101/6/BAMS-D-19-0015.1.xml>
- [8] R. Hine, S. Willcox, G. Hine, and T. Richardson, “The wave glider: A wave-powered autonomous marine vehicle,” in *OCEANS 2009*, 2009, pp. 1–6.
- [9] “SeaTrac.” [Online]. Available: <https://www.seatrac.com/>
- [10] A. Sekhar, B. Manoj, C. Siva, and R. Murthy, “Dynamic coverage maintenance algorithms for sensor networks with limited mobility,” in *Third IEEE International Conference on Pervasive Computing and Communications*. IEEE, 2005, pp. 51–60.
- [11] B. Liu, O. Dousse, P. Nain, and D. Towsley, “Dynamic coverage of mobile sensor networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 2, pp. 301–311, 2012.
- [12] C. Song, L. Liu, G. Feng, Y. Wang, and Q. Gao, “Persistent awareness coverage control for mobile sensor networks,” *Automatica*, vol. 49, no. 6, pp. 1867–1873, 2013.
- [13] A. Bafandeh and C. Vermillion, “Optimal altitude control of an integrated airborne wind energy system with globalized Lyapunov-based switched extremum seeking,” in *2018 European Control Conference (ECC)*. IEEE, 2018, pp. 2089–2094.
- [14] I. Bogunovic, J. Scarlett, and V. Cevher, “Time-varying Gaussian process bandit optimization,” in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 314–323.
- [15] A. Carpentier, A. Lazaric, M. Ghavamzadeh, R. Munos, and P. Auer, “Upper-confidence-bound algorithms for active learning in multi-armed bandits,” in *International Conference on Algorithmic Learning Theory*. Springer, 2011, pp. 189–203.
- [16] E. Contal, D. Buffoni, A. Robicquet, and N. Vayatis, “Parallel Gaussian process optimization with upper confidence bound and pure exploration,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2013, pp. 225–240.
- [17] E. Vazquez and J. Bect, “Convergence properties of the expected improvement algorithm with fixed mean and covariance functions,” *Journal of Statistical Planning and Inference*, vol. 140, no. 11, pp. 3088–3095, 2010.
- [18] G. A. Hollinger and G. S. Sukhatme, “Sampling-based motion planning for robotic information gathering,” in *Robotics: Science and Systems*, vol. 3, no. 5. Citeseer, 2013.
- [19] K. Cesare, R. Skeele, S.-H. Yoo, Y. Zhang, and G. Hollinger, “Multi-UAV exploration with limited communication and battery,” in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 2230–2235.
- [20] H. Ergezer and K. Leblebicioglu, “Path planning for UAVs for maximum information collection,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 502–520, 2013.
- [21] G. Zhang, S. Ferrari, and M. Qian, “An information roadmap method for robotic sensor path planning,” *Journal of Intelligent and Robotic Systems*, vol. 56, no. 1, pp. 69–98, 2009.
- [22] P. Hart, N. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [23] K. Chen and R. He, “Numerical investigation of the Middle Atlantic Bight shelfbreak frontal circulation using a high-resolution ocean hindcast model,” *Journal of Physical Oceanography*, vol. 40, no. 5, pp. 949–964, 2010.
- [24] D. P. Dee, S. M. Uppala, A. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. Balmaseda, G. Balsamo, d. P. Bauer, et al., “The era-interim reanalysis: Configuration and performance of the data assimilation system,” *Quarterly Journal of the royal meteorological society*, vol. 137, no. 656, pp. 553–597, 2011.
- [25] W. Bentz, T. Hoang, E. Bayasgalan, and D. Panagou, “Complete 3-D dynamic coverage in energy-constrained multi-UAV sensor networks,” *Autonomous Robotics*, vol. 42, pp. 825 – 851, 2018.