# Maximizing Seaweed Growth on Autonomous Farms: A Dynamic Programming Approach for Underactuated Systems Operating in Uncertain Ocean Currents

Matthias Killer[1,2,*], Marius Wiggert[1,*], Hanna Krasowski[2], Manan Doshi[3], Pierre F.J. Lermusiaux[3] and Claire J. Tomlin[1]

***Abstract*— Seaweed biomass presents a substantial opportunity for climate mitigation, yet to realize its potential, farming must be expanded to the expansive open oceans. However, in the open ocean neither anchored farming nor floating farms operating with powerful engines are economically viable. Recent studies have shown that vessels can navigate with low-power engines by *going with the flow*, utilizing minimal propulsion to strategically leverage beneficial ocean currents. In this work, we focus on low-power autonomous seaweed farms and design controllers that maximize seaweed growth by taking advantage of ocean currents. We first introduce a Dynamic Programming (DP) formulation to solve for the growth-optimal value function when the true currents are known. However, in reality only short-term imperfect forecasts with increasing uncertainty are available. Hence, we present three additional extensions. Firstly, we use frequent replanning to mitigate forecast errors. For that we compute the value function daily as new forecasts arrive, which also provides a feedback policy that is equivalent to replanning on the forecast at every time step. Second, to optimize for long-term growth, we extend the value function beyond the forecast horizon by estimating the expected future growth based on seasonal average currents. Lastly, we introduce a discounted finite-time DP formulation to account for the increasing uncertainty in future ocean current estimates. We empirically evaluate our approach with 30-day simulations of farms in realistic ocean conditions. Our method achieves 95.8% of the best possible growth using only 5-day forecasts. This confirms the feasibility of using low-power propulsion to operate autonomous farms in real-world conditions.***

## I. Introduction

Recent research has shown promising applications of seaweed biomass for climate mitigation. It can be used as human food, as cattle feed that reduces methane emissions [1], for biofuel and plastic [2], and for carbon capture i.e. when the biomass is sunk to the ocean floor, it removes carbon dioxide from the atmosphere [3]. To deliver on this promise, production must scale by expanding seaweed farming from labor-intensive operations near shore to automated solutions utilizing the vast expanse of the open oceans [4]. But conventional farming becomes economically infeasible in deeper waters as anchoring costs increase with depth [5].

* M.K. and M.W. have contributed equally to this work.

1 M.K., M.W., and C.J.T. are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA. For inquiries contact: `mariuswiggert@berkeley.edu`

2 M.K. and H.K. are with the Technical University of Munich, Germany

3 M.D. and P.F.J.L. are with the Department of Mechanical Engineering at the Massachusetts Institute of Technology, USA.

The authors gratefully acknowledge the support of the C3.ai Digital Transformation Institute and the IFI program of the German Academic Exchange Service (DAAD).

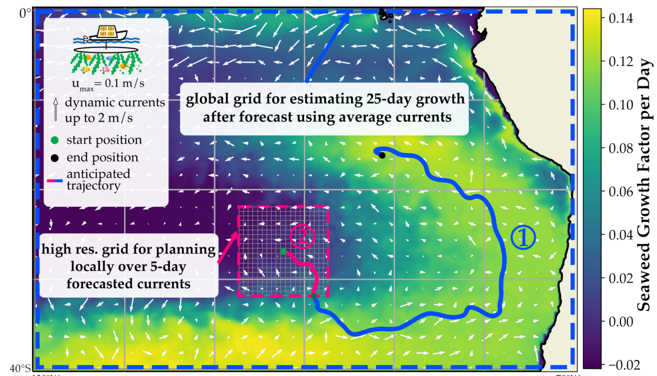


Solve for: $J^*_{Forecast}(\boldsymbol{x},t) = \int^{Forecast} growth(s)ds \; + \; \mathbb{E}\left[growth_{afterForecast}\right]$

Apply: $\pi_{Forecast}(\boldsymbol{x},t) = \arg\max_{\boldsymbol{u}\in\mathbb{U}} f(\boldsymbol{x},\boldsymbol{u},t)\cdot\nabla_x J^*_{Forecast}(\boldsymbol{x},t)$

Fig. 1: Our method maximizes long-term growth on autonomous seaweed farm that operate by harnessing ocean currents. We solve for the value function $J^*_{Forecast}$ that is long-term growth-optimal under the forecast with dynamic programming. We first compute the expected 25-day growth after the forecast based on historical average currents (**1**) and then use it to regularly solve for the value function over the next 5 days using daily current forecasts (**2**). Applying the induced policy $\pi_{Forecast}$ as feedback controller ensures high growth despite imperfect short-term forecasts.

A promising solution could be non-tethered, autonomous seaweed farms that roam the oceans while growing seaweed [6], [7]. These floating farms needs to be able to control their position to prevent stranding, colliding with ships, or drifting to nutrient-depleted waters. While they could be steered with powerful ship engines, the power and hence energy costs are prohibitively high due to the drag force scaling quadratically with the relative velocity of the farm. In our recent work, we demonstrated that an Autonomous Surface Vehicle (ASV) can navigate reliably by *going with the flow*, using its minimal propulsion ($0.1\frac{m}{s}$) strategically to nudge itself into ocean currents ($[0-2\frac{m}{s}]$) that drift towards its destination [8], [9]. This work has been extended to reduce the risk of stranding [10] and to fleets of vessels that navigate while staying connected in a local communication network [11]. In this paper, we use this low-power steering paradigm for operating seaweed farms. Our objective is to maximize seaweed growth along the trajectory of the farms building upon prior research on optimal deterministic autonomous sea farming [12], [13]. From the control perspective, there are four key challenges that we need to tackle. First, the currents are non-linear and time-varying. Second, in realistic settings, only coarse uncertain forecasts are available [14]–

[18]. Third, the farm itself is *underactuated* by which we mean that its propulsion is smaller than the surrounding currents, so it cannot easily compensate for forecast errors. Lastly, we want to maximize seaweed growth over weeks but forecasts from the leading providers are only 5-10 days long [15], [16] and the uncertainty for long-time predictions is high [19], [20]. In a nutshell, we are tackling long-term horizon optimization of a state-dependent running cost with an underactuated agent in non-linear time-varying dynamics under uncertainty that increases over time. The long-term dependency of seaweed growth means the objective cannot easily be decomposed into multiple short-term objectives.

### A. Related Work

Various approaches for time- and energy-optimal path planning exist for non-linear, time-varying dynamics like ocean currents [21]–[37]. In the context of navigating within *known currents or flows*, researchers have derived Hamilton-Jacobi (HJ) reachability equations for exact solutions [21]–[23], non-linear programming [24], [38], evolutionary algorithms [25], and graph-based search methods [26], [28], [36], [39]. However, the last three techniques are prone to discretization errors and the non-convex nature of the problem, can lead to infeasibility or solvers getting stuck in local minima. In contrast, DP based on the HJ equations can solve the exact continuous-time control problem.

There is only a little research that focuses on maximizing seaweed growth. In [12], the authors maximize seaweed harvesting using autonomous vessels in varied settings. They use a 3D HJ reachability framework in which the harvesting state is augmented into the third dimension. To find the path with maximum growth, they run forward reachability in the state space for seaweed in 3D. This formulation needs to be adapted for closed-loop control and for realistic operational conditions accounting for ocean forecast uncertainties.

For managing current uncertainty, previous work optimizes the expectation or a risk-function over a stochastic solution of probabilistic ocean flows [40]. However, this is not yet suitable for operational settings as it demands a principled uncertainty distribution for flows but most operational forecasts are deterministic. At the same time, robust control techniques, which aim to maximize the objective even in the face of worst-case disturbances, are not suitable when considering realistic error bounds, as the forecast error often equals or exceeds our low propulsion capabilities. Thus, to mitigate forecast inaccuracies, frequent replanning in a Model Predictive Control (MPC) fashion has been proposed using either non-linear programming [41], [42] or employing the HJ value function as feedback policy [8], which offers the benefits of being both fast and optimal. An emergent approach is to use Reinforcement Learning (RL) to learn how to best operate stratospheric balloons despite wind forecast uncertainty [37], [43]. RL is suitable for their short-term objective of station keeping and it is unclear if this works for long-term objectives such as ours.

In order to address the increasing complexity associated with long-time horizons, problems are frequently divided into multiple subproblems using graph-based methods or hierarchical RL [44], [45]. These approaches are appropriate for combinatorial optimization problems, where dividing and conquering in subtasks is effective. However, this is not suitable for our problem involving continuous state space and long-time dependencies. A potential solution to handle growing uncertainty over time is to discount future rewards. This is common in RL settings [46], [47].

### B. Overview of Method & Contributions

In this paper, we make five main contributions towards controllers that optimize seaweed growth on autonomous seaweed farms over long periods.

First, we formulate maximizing a type of seaweed growth as an optimization problem that can be solved exactly with DP in the 2D spatial state of the system (Sec. III-A). Compared to prior work using HJ reachability in 3D [12] our assumptions lead to two advantages: significant reduction of computational complexity; value function that can be used as feedback policy to obtain the growth-optimal control for all states and times, allowing frequent replanning in MPC spirit for multiple farms which is critical when only forecasts are known. Second, we extend the value function beyond the forecast horizon which leads to a feedback policy that optimizes for long-term optimal growth (Sec. III-C). Third, to account for the growing uncertainty of the ocean current estimates, we introduce finite-time discounting into the DP formulation (Sec. III-D). Forth, we are the first to run extensive empirical simulations of autonomous seaweed farms in realistic current settings over 30 days. We first investigate how different propulsion of the farms would affect the best achievable seaweed growth with known currents. We then evaluate how close different configurations of our method can get to the best achievable growth when only daily, 5-day forecasts are available (Sec. IV). Lastly, we open-source our code-base which contains extensive features to simulate, visualize, and study controllers for 2D vessels operating by harnessing uncertain ocean currents.

In Sec. II we define the problem. Sec. III details the four components of our method. Sec. IV contains the performance evaluation of our methods and baselines and we conclude with Sec. VI and outline future work.

## II. Problem Statement

### A. System Dynamics

We consider an autonomous seaweed farm as surface vessel on the ocean with the spatial state $\boldsymbol{x} \in \mathbb{R}^2$. Let the control input be denoted by $\boldsymbol{u}$ from a bounded set $\mathbb{U} \in R^{n_u}$ where $n_u$ is the dimensionality of the control. Then, the spatial dynamics of the system at time $t$ can be modelled by the first order Ordinary Differential Equation (ODE):

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}, \boldsymbol{u}, t) = v(\boldsymbol{x}, t) + g(\boldsymbol{x}, \boldsymbol{u}, t), \quad t \in [0, T] \qquad (1)$$

where the movement of the vessel depends on the drift due to the time-varying, non-linear flow field $v(\boldsymbol{x}, t) \rightarrow \mathbb{R}^2$ and its control $\boldsymbol{u}$. We choose a first-order model where the drift and control directly influence the state, disregarding inertial effects from motor acceleration and drag forces. This is justified by the fact that high-drag seaweed farms attain equilibrium velocity within a few minutes, a timescale considerably shorter than our 30-day planning horizon.

While our method is generally applicable, we focus on *underactuated* settings in the sense that most of the time $\max_{\boldsymbol{u}} \|g(\boldsymbol{x},\boldsymbol{u},t)\|_2 \ll \|v(\boldsymbol{x},t)\|_2$. We denote the spatial trajectory induced by this ODE with $\boldsymbol{\xi}$. For a vessel starting at the initial state $\boldsymbol{x}_0$ at time $t_0$ with control sequence $\boldsymbol{u}(\cdot)$, we denote the state at time $t$ by $\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t_0,\boldsymbol{x}_0}(t) \in \mathbb{R}^2$. The system dynamics (Eq. 1) are assumed to be continuous, bounded, and Lipschitz continuous in $\boldsymbol{x}, \boldsymbol{u}$ [9].

Additionally, we assume the farm has seaweed mass $m$ which evolves according an exponential growth ODE:

$$\dot{m} = m \cdot \Psi(\boldsymbol{x},t), \quad t \in [0,T] \tag{2}$$

where $\Psi$ is the growth factor per time unit, e.g. $20\frac{\%}{\text{day}}$, which depends on nutrients, incoming solar radiation, and water temperature at the spatial state $\boldsymbol{x}$ and time $t$.

### B. Problem Setting

The objective of the seaweed farm starting from $\boldsymbol{x}_0$ at $t_0$ with seaweed mass $m(t_0)$ is to maximize the seaweed mass at the final time $T$. This implies optimizing the growth over its spatial trajectory $\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t_0,\boldsymbol{x}_0}$.

$$\max_{\boldsymbol{u}(\cdot)} m(T) = m(t_0) + \max_{\boldsymbol{u}(\cdot)} \int_{t_0}^{T} m(s) \cdot \underbrace{\Psi(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t_0,\boldsymbol{x}_0}(s), s)}_{\text{growth factor}} \, ds \tag{3}$$

If the currents $v$ are known, our method (Sec. III) is guaranteed to find the optimal value function from which the optimal control $\boldsymbol{u}^*(\cdot)$ and trajectory can be obtained. However, in realistic scenarios only inaccurate, short-term forecasts $\hat{v}_{FC}$ are available at regular intervals. These differ from the true flow $v$ by the forecast error $\delta(\boldsymbol{x},t)$. Our goal is then to determine a feedback policy $\pi(\boldsymbol{x},t)$ that results in a high expected seaweed mass $\mathbb{E}[m(T)]$. Hence, in our experiments (Sec. IV) we evaluate our method empirically over a set of missions $(\boldsymbol{x}_0,t) \sim \mathbb{M}$ and a realistic distribution of true and forecasted ocean currents $v, \hat{v}_{FC} \sim \mathbb{V}$.

## III. Method

Our method consists of a core DP formulation that optimizes seaweed growth when the currents are known and three extensions to get a feedback policy $\pi$ that performs well over long-time horizons when only forecasts are available. We first introduce the core DP formulation to obtain the growth-optimal value function (Sec. III-B). Then we show how it can be used as feedback policy $\pi$ that is equivalent to replanning at every time step (Sec. III-C). Next, we show how this can be extended beyond the forecast horizon (Sec. III-C). Lastly, we introduce a finite-time discount factor in the DP formulation (Sec. III-D).

### A. Maximizing Seaweed Mass With Known Dynamics

We use continuous-time optimal control where the value function $J(\boldsymbol{x},\boldsymbol{u}(\cdot),t)$ of a trajectory $\boldsymbol{\xi}$ is based on a state and time-dependent reward $R$ and a terminal reward $R_T$:

$$J(\boldsymbol{x},\boldsymbol{u}(\cdot),t) = \int_t^T R(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(s), s) ds + R_T(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(T), T).$$

Let $J^*(\boldsymbol{x},t) = \max_{\boldsymbol{u}(\cdot)} J(\boldsymbol{x},\boldsymbol{u}(\cdot),t)$ be the optimal value function. Using DP we can derive the corresponding Hamilton-Jacobi Partial Differential Equation (PDE) [48]:

$$-\frac{\partial J^*(\boldsymbol{x},t)}{\partial t} = \max_{\boldsymbol{u}} \left[\nabla_{\boldsymbol{x}} J^*(\boldsymbol{x},t) \cdot f(\boldsymbol{x},\boldsymbol{u},t) + R(\boldsymbol{x},t)\right] \tag{4}$$

$$J^*(\boldsymbol{x},T) = R_T(\boldsymbol{x},T). \tag{5}$$

We can then numerically compute $J^*(\boldsymbol{x},t)$ on a spatial mesh by integrating the PDE backwards in time [49].

Next, we define the reward $R$ and terminal reward $R_T$ to maximize $m(T)$. One approach is to solve the PDE in an augmented state space $x_{aug} = (\boldsymbol{x}, m)^\top \in \mathbb{R}^3$. If we set $R_T = 0$ and define the reward as $R = m \cdot \Psi(\boldsymbol{x},t)$, the value function is our objective (Eq. 3). However, the computational complexity of solving for $J^*$ scales rapidly with the state dimension. Hence, we want a reward $R$ that does not depend on $m$ as augmented state. For that, we introduce the variable $\eta = \ln(m)$ with the new dynamics $\dot{\eta} = \frac{\dot{m}}{m} = \Psi(\boldsymbol{x},t)$. As $\eta(m)$ is strictly increasing in $m$, the control $\boldsymbol{u}^*(\cdot)$ that maximizes $\eta(T)$ is equivalent to $\boldsymbol{u}^*(\cdot)$ maximizing $m(T)$. We can then reformulate Eq. 3 to $\eta(T)$:

$$\max_{\boldsymbol{u}(\cdot)} \eta(T) = \eta(t_0) + \max_{\boldsymbol{u}(\cdot)} \int_{t_0}^{T} \Psi(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(s), s) \, ds. \tag{6}$$

By setting the reward to $R = \Psi(\boldsymbol{x},t)$ the optimal value function captures this optimization without requiring $m$:

$$J^*(\boldsymbol{x},t) = \max_{\boldsymbol{u}(\cdot)} \int_{t}^{T} \Psi(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(s), s) \, ds. \tag{7}$$

We then solve the HJ PDE for the growth-optimal $J^*(\boldsymbol{x},t)$ in the spatial state $\boldsymbol{x}$ and obtain $\boldsymbol{u}^*(\cdot)$ and the trajectory $\boldsymbol{\xi}^{\boldsymbol{u}^*(\cdot)}_{t_0,\boldsymbol{x}_0}$ that maximize $m(T)$ at lower computational cost. We can convert the value of $J^*(\boldsymbol{x}_0,t_0)$ to the final seaweed mass of the optimal trajectory starting at $\boldsymbol{x}_0, t_0$ with $m(t_0)$:

$$m(T) = m(t_0) \cdot e^{\int_{t_0}^{T} \Psi(\boldsymbol{\xi}^{\boldsymbol{u}^*(\cdot)}_{t_0,\boldsymbol{x}_0}(s),s)\, ds} = m(t_0) \cdot e^{J^*(\boldsymbol{x}_0,t_0)}.$$

### B. Feedback Policy Based on Regular Forecasts

The value function $J^*$ from Sec. III-A allows us to compute the optimal control $\boldsymbol{u}^*(\boldsymbol{x},t)$ for all $\boldsymbol{x}, t$ and hence a feedback policy $\pi(\boldsymbol{x},t)$ for the vessel or multiple vessels in the same region [8]. This policy is the optimizer of the Hamiltonian (right side Eq. 4):

$$\pi(\boldsymbol{x},t) = \arg\max_{\boldsymbol{u}\in\mathbb{U}} f(\boldsymbol{x},\boldsymbol{u},t) \cdot \nabla_x J^*(\boldsymbol{x},t), \tag{8}$$

which can often be computed analytically depending on $g(\boldsymbol{x},\boldsymbol{u},t)$. While $\pi$ is optimal if $J^*$ is based on the true currents $v$, it can also be applied when imperfect forecasts $\hat{v}_{FC}$ were used to compute the value function $J^*_{\hat{v}_{FC}}(\boldsymbol{x},t)$. In that case, an agent at state $\boldsymbol{x}$ executing $\pi_{\hat{v}_{FC}}(\boldsymbol{x},t)$ will find itself at a different state $\boldsymbol{x}'$ than anticipated as $v$ differs from $\hat{v}_{FC}$. But the control that would be growth optimal under $\hat{v}_{FC}$ can again be computed with $\pi_{\hat{v}_{FC}}(\boldsymbol{x}', t+\Delta t)$. Applying $\pi_{\hat{v}_{FC}}$ closed-loop is hence equivalent to full-time horizon re-planning with $\hat{v}_{FC}$ at each time step. This notion of re-planning at every time step at a low computational cost ensures good performance despite forecast errors [8]. $J^*_{\hat{v}_{FC}}(\boldsymbol{x},t)$ can be updated daily when new forecasts arrive.

### C. Reasoning Beyond the Forecast Horizon

As the growth cycles of seaweed typically spans months, our aim is to maximize the seaweed mass at an *extended* future time $T_{\text{ext}}$ after the final time of the 5-day forecast $T_{FC}$. A principled way to reason beyond the planning horizon is to estimate the expected growth our seaweed farm will experience from the state $\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(T_{FC})$ onward and add this as terminal reward $R_T$ to Eq. 7.

$$J^*_{\hat{v}_{FC},\text{ext}}(\boldsymbol{x},t) = J^*_{\hat{v}_{FC},T_{FC}}(\boldsymbol{x},t) + \mathbb{E}\left[J^*_{T_{\text{ext}}}(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(T_{FC}),T_{FC})\right]$$

$$J^*_{\hat{v}_{FC},T_{FC}}(\boldsymbol{x},t) = \max_{\boldsymbol{u}(\cdot)} \int_t^{T_{FC}} \Psi(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(s),s)\, ds \quad (9)$$

where $J^*_{\hat{v}_{FC},T_{FC}}(\boldsymbol{x},t)$ is the growth a vessel starting from $\boldsymbol{x}$ at $t$ will achieve at $T_{FC}$ and $\mathbb{E}\left[J^*_{T_{\text{ext}}}(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(T_{FC}),T_{FC})\right]$ estimates the additional growth from $T_{FC}$ to $T_{\text{ext}}$. The expectation is over the uncertain future ocean currents.

We propose to estimate $\mathbb{E}\left[J^*_{T_{\text{ext}}}\right]$ by computing the value function $J^*_{\bar{v},T_{\text{ext}}}$ based on monthly average currents $\bar{v}$ for the region using Sec. III-A. To compute $J^*_{\hat{v}_{FC},\text{ext}}$ we then solve Eq. 4 with $R_T(\boldsymbol{x},T_{FC}) = J^*_{\bar{v},T_{\text{ext}}}(\boldsymbol{x},T_{FC})$.

### D. Finite-time Discounting to Mitigate Uncertainty

As the oceans are a chaotic system, the uncertainty of the forecasted ocean currents increases over time. We can incorporate this increasing uncertainty in the value function by using the finite-time discounted optimal control formulation:

$$J^{\tau}(\boldsymbol{x},\boldsymbol{u}(\cdot),t) = \int_t^T e^{\frac{-(s-t)}{\tau}} R(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(s),s)\, ds + R_T(\boldsymbol{\xi}^{\boldsymbol{u}(\cdot)}_{t,\boldsymbol{x}}(T),T),$$

where $\tau$ is the discount factor. The smaller $\tau$ the more future rewards are discounted. We derive the corresponding HJ PDE by following the steps in [47] and in place of Eq. 4 we obtain:

$$\frac{\partial J^{*,\tau}(\boldsymbol{x},t)}{\partial t} = -\max_{\boldsymbol{u}}\left[\nabla_{\boldsymbol{x}} J^{*,\tau}\cdot f(\boldsymbol{x},\boldsymbol{u},t) + R(\boldsymbol{x},t)\right] + \frac{J^{*,\tau}(\boldsymbol{x},t)}{\tau}$$

### E. Summary Control Algorithm Variations

All variations of our method are feedback policies $\pi$ derived from a value function (Sec. III-B). The four variations differ only in how the value function is computed. When the true currents $v$ are known we compute $J^*$ (Eq. 7) for optimal control. When only forecasts $\hat{v}_{FC}$ are available, we calculate the $J^*_{\hat{v}_{FC}}$ for planning horizons up to the end of the forecasts $T_{FC}$ and update it as new forecasts become available (Sec. III-B). Thirdly, to optimize for growth until $T_{\text{ext}} > T_{FC}$ we calculate an extended value function $J^*_{\hat{v}_{FC},\text{ext}}$ (Sec. III-C) using average currents $(\hat{v}_{FC} + \bar{v})$. Lastly, we can discount future rewards with $J^{*,\tau}$ (Sec. III-D) in any of the above value functions. In Algorithm 1 we detail the discounted, long-term version as it contains all components.

## IV. Experiments

In this section, we empirically evaluate various settings of our method for operating an autonomous seaweed farm in realistic ocean currents and growth conditions over $T$=30 days. Our farm has actuation $g(\boldsymbol{x},\boldsymbol{u},t) = \boldsymbol{u}$ with varying $\boldsymbol{u}_{max}$. We open-sourced the code for our simulator and controllers for others to replicate results and build on [1].

**Algorithm 1:** Discounted HJ Closed-loop Control

**Input:** Forecast Flow(s) $\hat{v}_{FC}$, $t=0$, $\boldsymbol{x}(t)=\boldsymbol{x}_0$, average Flows $\bar{v}$, discount $\tau$, plan until $T_{\text{ext}}$

```
1 Compute J^{*,τ}_{v̄,T_ext} using v̄ (Sec. III-C);
2 while t ≤ T do
3     if new forecast v̂_FC available then
4         Compute J^{*,τ}_{v̂_FC,ext} (Sec. III-C);
5     u_t = π^{*,τ}_{v̂_FC,ext}(x_t, t); using J^{*,τ}_{v̂_FC,ext} (Sec. III-B)
6     x(t + Δt) = x(t) + ∫_t^{t+Δt} f(u_t, x(s), s) ds;
7     t ← t + Δt;
```

We run two experiments. First we investigate how varying the propulsion $\boldsymbol{u}_{max}$ impacts the best achievable seaweed growth under known currents $v$ and compare it to the growth achieved by 30-day planning without discounting relying on daily, 5-day forecasts $v$ and average currents $\bar{v}$ (Sec. IV-B.1). Second, we fix the propulsion to $\boldsymbol{u}_{max} = 0.1\frac{m}{s}$ and evaluate how the planning horizon and discounting in our method affect growth and how close we can get to the best achievable growth while relying on daily forecasts $v$ and average currents $\bar{v}$ (Sec. IV-B.2). The experimental setup for both is the same and will be explained next.

### A. Experimental Setup

*1) Seaweed Growth Model:* Macroalgae growth depends on the species, the water temperature, solar irradiance, and dissolved nutrient concentrations, specifically nitrate $(\mathrm{NO_3})$ and phosphate $(\mathrm{PO_4})$ [12]. We use the model of the Net Growth Rate (NGR) of Wu et al. [50] and temperate species parameters from [51], [52]. In this model, the time-dependent NGR is determined by the growth rate $r_{growth}$ and the respiration rate $r_{resp}$ caused by metabolism as:

$$\dot{m}(t) = m(t)\cdot \mathrm{NGR}(t) = m(t)\cdot(r_{growth}(t) - r_{resp}(t)). \quad (10)$$

Fig. 1 shows the NGR for our region at the apex of the sun's motion in January 2022.

*2) Realistic Ocean Forecast Simulation:* In realistic operations the vessel receives daily forecasts for replanning. In our simulations, we use Copernicus [16] hindcasts as true currents $v$ and mimic daily 5-day forecasts $\hat{v}_{FC}$ by giving the planner access to a 5-day sliding time window of HYCOM [53] hindcasts. As in our previous work [54] we find that the forecast error $\delta$ with this setting is comparable to the evaluated forecast error of HYCOM [14] in key metrics. To estimate the expected growth beyond the forecast horizon of $\hat{v}_{FC}$ (Sec. III-C) we use $\frac{1}{6}$th deg seasonal averages $\bar{v}$ of the ocean currents from Copernicus 2021.

*3) Large Scale Mission Generation:* We simulate operations in the southeast Pacific due to high nutrient densities. For a large representative set of missions $\mathbb{M}$, we sampled 1325 tuples $(\boldsymbol{x}_0, t_0, m(t_0) = 100\text{kg})$, uniformly distributed in time between January and October 2022 and across the

[1] https://github.com/MariusWiggert/OceanPlatformControl

| controller | planning horizon $T_{\text{ext}}$ | discount factor $\tau$ |
|---|---|---|
| w/o discount ($v$) | 30 days | - |
| floating | - | - |
| greedy 1 hour ($\hat{v}_{FC}$) | 1 hour | - |
| greedy 5 days ($\hat{v}_{FC}$) | 5 days | - |
| w/o discount ($\hat{v}_{FC}$ + $\bar{v}$) | 30 days | - |
| w/ discount I ($\hat{v}_{FC}$ + $\bar{v}$) | 30 days | 1.296.000 |
| w/ discount II ($\hat{v}_{FC}$ + $\bar{v}$) | 30 days | 1.728.000 |

TABLE I: We compare various controller settings.

region of longitude [-130, -70]°W and latitude [-40, 0]°S. This allows for varying current distributions. As our method is not aware of land obstacles we had 290 missions where at least one of the controllers stranded or left the simulation region. While stranding can be avoided by modifying the HJ PDE as demonstrated in parallel work [10], we consider only the remaining 1035 missions for our results.

*4) Evaluated Controllers:* We evaluate our method in different configurations categorized by: a) the ocean current data used by the controller for planning, either the true currents $v$ or daily forecasts $\hat{v}_{FC}$ and average currents $\bar{v}$, and b) the controller's planning horizon $T_{\text{ext}}$ over which it optimizes growth, either the entire 30-day period or more short term *greedy* (5-day and 1 hour). We also examine the use of a discounted value function. We compare all controllers against the scenario where the farms float passively. An overview of the configurations is provided in Tab. I. For long-term ($T_{\text{ext}}$=30 days) controllers, we compute the growth-to-go after $T_{FC}$, $J^*_{\bar{v},T_{\text{ext}}}(\boldsymbol{x}, T_{FC})$, over the full area on a coarse $\frac{1}{6}°$ grid, as illustrated in Fig. 1. The value function $J^*_{\hat{v}_{FC},\text{ext}}(\boldsymbol{x}, t)$ used for the control policy is then computed daily on new forecasts using a smaller $\frac{1}{12}°$ grid around the current farm's position (10° square).

*5) Evaluation Metrics:* Our objective is to maximize the seaweed mass at the end of each mission $m(T)$. Additionally, we compute the relative improvement in final seaweed mass by normalizing within each mission with the baseline final mass. We then present the average relative improvement across all missions which allows us to gauge how much more/less biomass a specific controller can grow above the baseline. This is important as the start $\boldsymbol{x}_0$ of a mission is a major indicator of achievable growth as illustrated in Fig. 3. As baselines we use either passively floating or the best achievable growth based on the true currents $v$.

### B. Experimental Results

*1) How does varying propulsion affect growth?:* We vary the maximum propulsion $\boldsymbol{u}_{max}$ of the farm and evaluate how this impacts the best achievable seaweed growth under known currents $v$. Fig. 2 and Tab. II compare the final seaweed mass distributions for different propulsion levels, starting with passively floating. We observe that the average seaweed growth scales almost linearly with $\boldsymbol{u}_{max}$, yielding between 15% and 12% more biomass per $0.1\frac{m}{s}$ propulsion. We also compare how much growth our method w/o discount ($\hat{v}_{FC}$ + $\bar{v}$) achieves with varying propulsion. As expected this achieves slightly less biomass (≈95-96% of $v$) due to forecast errors for all propulsion levels. For higher $\boldsymbol{u}_{max}$

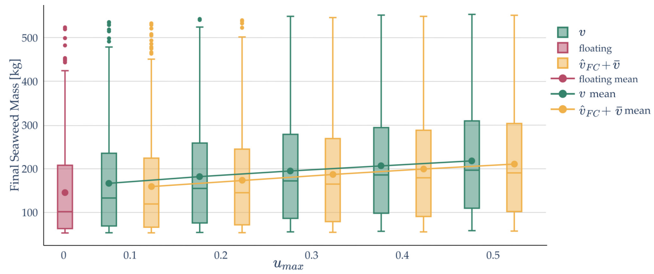


Fig. 2: The best achievable seaweed mass given $v$ increases linearly with $u_{max}$. Operating with our long-term control method using forecasts $\hat{v}_{FC}$ and average currents $\bar{v}$ achieves ≈ 95% of growth.

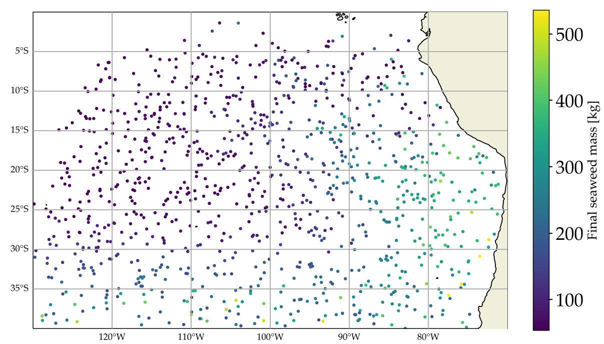


Fig. 3: We sample a diverse set of starts $(\boldsymbol{x}_0, t_0)$ for seaweed farms to empirically evaluate our controllers. The colorized starts show the best achievable seaweed mass after 30 days using $u_{max} = 0.1\frac{m}{s}$.

the gap is slightly smaller, possibly because the farm can better compensate for forecast errors. Nonetheless, even small propulsion of $u_{max}$=$0.1\frac{m}{s}$ enables 9.6% more biomass than a passively floating farm.

The start $\boldsymbol{x}_0$ of a mission significantly influences 30-day growth, as shown in Fig. 3. High-growth missions are situated in the east and south of our region, aligning with nutrient-rich areas (see Fig. 1).

| $u_{max}$ | planning input | rel. growth | final seaweed mass |
|---|---|---|---|
| $0.0\frac{m}{s}$ | (floating) | 100% | 145.29kg±100.30kg |
| $0.1\frac{m}{s}$ | $v$ | 115.38% | 166.45kg±109.67kg |
| | $\hat{v}_{FC}$ + $\bar{v}$ | 109.62% | 159.29kg±107.46kg |
| $0.2\frac{m}{s}$ | $v$ | 128.69% | 182.04kg±115.11kg |
| | $\hat{v}_{FC}$ + $\bar{v}$ | 121.29% | 173.72kg±112.94 |
| $0.3\frac{m}{s}$ | $v$ | 141.27% | 194.98kg±117.39kg |
| | $\hat{v}_{FC}$ + $\bar{v}$ | 133.28% | 187.01kg±116.60kg |
| $0.4\frac{m}{s}$ | $v$ | 153.71% | 206.96kg±118.34kg |
| | $\hat{v}_{FC}$ + $\bar{v}$ | 145.79% | 199.50kg±118.09kg |
| $0.5\frac{m}{s}$ | $v$ | 165.79% | 218.10kg±118.59kg |
| | $\hat{v}_{FC}$ + $\bar{v}$ | 158.14% | 210.78kg±117.72kg |

TABLE II: Average seaweed growth for different maximum propulsions $\boldsymbol{u}_{max}$. Planning on $v$ is the best achievable which we compare to using forecasted and average currents ($\hat{v}_{FC}$ + $\bar{v}$) without discounting. Normalized per mission by passively floating.

*2) The impact of planning horizon and discounting:* As the energy consumption scales cubically with $\boldsymbol{u}_{max}$, higher propulsion may be economically infeasible for real-world applications. Therefore, for this experiment we fix $\boldsymbol{u}_{max}$= $0.1\frac{m}{s}$. We investigate how different planning horizons and discounting affect performance when operating with fore-

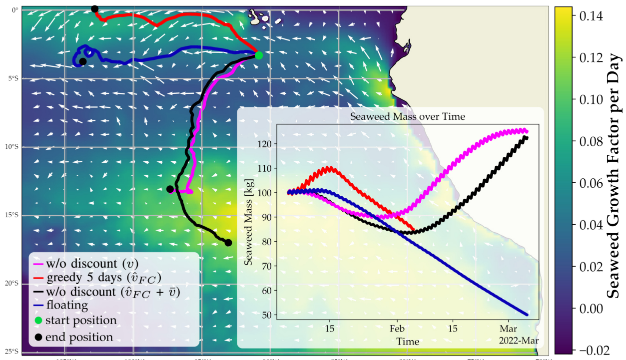


Fig. 4: 60-day Case Study: The greedy controller optimizes for 5-day growth thereby navigating to the closest growth region. It fails to anticipate the strong currents that push it out of the region. The long-term controllers reach a more distant growth-richer area while incurring short-term losses.

casts $\hat{v}_{FC}$ and how close we can get to the best achievable growth. We evaluate two greedy controllers that repeatedly optimize over short $T_{\text{ext}}$ (1h and 5-days) and compare to 30-day time-horizon with different discounting settings (Tab. I).

Table III shows the results. As expected, both the greedy and long-term controllers outperform passively floating. Surprisingly, the performance of the 5-day greedy controller, is close to that 30-day controllers. Using the discounted formulation slightly improves the long-term controller, yielding on average 95.77% of the best achievable growth.

| **controller $u_{max} = 0.1\frac{m}{s}$** | **relative growth** | **final seaweed mass** |
|---|---|---|
| w/o discount ($v$) | 100% | 168.45kg±109.67kg |
| floating (-) | 88.20% | 145.29kg±99.54kg |
| greedy 1 hour ($\hat{v}_{FC}$) | 92.24% | 152.48kg±102.89kg |
| greedy 5 days ($\hat{v}_{FC}$) | 95.19% | 157.78kg±106.04kg |
| w/o discount ($\hat{v}_{FC}$ + $\bar{v}$) | 95.61% | 158.84kg±106.71kg |
| w/ discount I ($\hat{v}_{FC}$ + $\bar{v}$) | 95.77% | 159.16kg±106.62kg |
| **w/ discount II ($\hat{v}_{FC}$ + $\bar{v}$)** | **95.77%** | **159.17kg**±106.66kg |

TABLE III: Average seaweed growth of different controllers over 1035 missions. Normalized per mission by best achievable given $v$.

*3) Case Study of 60-Day Scenario:* We were intrigued that the 5-day controller did achieve almost the same seaweed growth as by planning over 30-days (Sec. IV-B.2). Hence, we conducted a case study with planning and operating the farm over 60 days instead of 30 days (Fig. 4). We find that the greedy controller then aims for the nearest growth region, while the long-term controller properly balances short-term losses against the long-term gains of reaching a high-growth region. This leads to the greedy controller being driven out of the region while the long-term controller achieves close to the best achievable growth (see sub-figure Fig. 4). The zig-zags shape of the lines are due to day-night cycles.

## V. Discussion

As expected, controllers using forecasts $\hat{v}_{FC}$ substantially outperform a passively floating farm. Since seaweed growth cycles span 60-90 days, we believe that long-term planning is crucial. The myopic behavior of a greedy policy not only leads it to navigate toward low growth regions in the vicinity but also fails to account for the possibility of being pushed out of good growth regions by strong currents, as in our 60-day case study in Fig. 4. Therefore, we were surprised that our 5-day optimizing controller was nearly on par with our 30-day optimizing controllers (Sect. IV-B.2).

We attribute this to several factors. First, the growth map in our region exhibits a smooth gradient, which means that even greedy controllers might move toward globally optimal growth regions without planning for it. Second, in our experimental evaluation, we do not consider missions where any controller leaves the predefined region (Sec. IV-A.3). This often occurs with greedy or floating controllers (Fig. 4); consequently, the performance increase with long-term controllers would be greater if we accounted for the filtered missions. Third, the 30-day simulation time-frame may not be enough to see the benefit of long-term planning as the farm can only travel a limited distance within that time. An indicator of this is the high variance in final seaweed mass, which can be attributed to the inability of our vessels to reach optimal growth regions in 30 days for many missions. Hence, we anticipate that for longer simulation times and higher propulsion $u_{max}$ we would see a higher performance divergence and reduced variance for long-term controllers.

## VI. Conclusion and Future Work

In this work, we maximize seaweed growth on autonomous farms that are underactuated and operate by harnessing uncertain ocean currents. We first introduced a 2D DP formulation to solve for the growth-optimal value function when the true currents are known. Next, we showed how the value function computed on forecasted currents can be used as feedback policy for multiple farms, which is equivalent to replanning on the forecast at every time step and hence mitigates forecast errors. As operational forecasts are only 5 days long, we extended our method to reason beyond the forecast horizon by estimating expected future growth based on seasonal average currents. Lastly, we presented a finite-time discounting DP PDE to account for increasing uncertainty in ocean currents. We conducted extensive empirical evaluations based on realistic ocean conditions over 30 days. Our method achieved 95.8% of the best achievable growth and 9.6% more growth than passively floating despite its low propulsion of $u_{max} = 0.1\frac{m}{s}$ and relying on daily 5-day forecasts. This confirms the feasibility of harnessing ocean currents to operate low propulsion autonomous seaweed farms.

A future direction is to learn the expected growth after the forecast horizon using experience and approximate value iteration [55] or a value network [56]. This could implicitly learn the distribution shift between $\hat{v}_{FC}$ and $v$. Another direction is to make the discount factor state-dependent based on the uncertainty of current predictions, which could be estimated historically or forecast [57], [58]. Further, we want to run experiments for time horizons longer than 30-days to tease out the advantage of long-term planning. Lastly, we plan to conduct field tests with our partner [6] to further validate our method in real-world ocean conditions.

## References


[1] B. M. Roque *et al.*, "Effect of the macroalgae asparagopsis taxiformis on methane production and rumen microbiome assemblage," *Animal Microbiome*, vol. 1, pp. 1–14, 2019.

[2] W. T. L. Yong *et al.*, "Seaweed: a potential climate change solution," *Renewable and Sustainable Energy Reviews*, vol. 159, no. 112222, 2022.

[3] "Seaweed farming as climate solution," Project Drawdown. [Online]. Available: https://drawdown.org/solutions/seaweed-farming

[4] R. M. Tullberg *et al.*, "Review of the status and developments in seaweed farming infrastructure," *Journal of Marine Science and Engineering*, vol. 10, no. 10, 2022.

[5] F. Ross *et al.*, "Seaweed afforestation at large-scales exclusively for carbon sequestration: Critical assessment of risks, viability and the state of knowledge," *Frontiers in Marine Science*, vol. 9, p. 2269, 2022.

[6] "Floating seaweed farms," Phykos PBC. [Online]. Available: https://www.phykos.co

[7] M. T. Sherman *et al.*, "Seaweedpaddock: Initial modeling and design for a sargassum ranch," in *OCEANS 2018 MTS/IEEE*, 2018, pp. 1–6.

[8] M. Wiggert *et al.*, "Navigating underactuated agents by hitchhiking forecast flows," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2417–2424.

[9] M. Doshi *et al.*, "Hamilton-jacobi multi-time reachability," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2443–2450.

[10] A. Doering *et al.*, "Stranding risk for underactuated vessels in complex ocean currents: Analysis and controllers," *arXiv preprint arXiv:2307.01917*, 2023.

[11] N. Hoischen *et al.*, "Safe connectivity maintenance in underactuated multi-agent networks for dynamic oceanic environments," *arXiv preprint arXiv:2307.01927*, 2023.

[12] M. S. Bhabra *et al.*, "Optimal harvesting with autonomous tow vessels for offshore macroalgae farming," in *OCEANS 2020 IEEE/MTS*. IEEE, Oct. 2020, pp. 1–10.

[13] M. S. Bhabra, "Harvest-time optimal path planning in dynamic flows," Master's thesis, Massachusetts Institute of Technology, Department of Mechanical Engineering and Computational Science & Engineering, Cambridge, Massachusetts, Sep. 2021.

[14] E. Metzger *et al.*, "Validation test report for the global ocean forecast system 3.5-1/25 degree hycom/cice with tides," Vaval Research Lab Washington DC, US, Tech. Rep., 2020.

[15] E. P. Chassignet *et al.*, "US GODAE: global ocean prediction with the hybrid coordinate ocean model (HYCOM)," *Oceanography*, vol. 22, no. 2, pp. 64–75, 2009.

[16] EU Copernicus Marine Service Information, "Global ocean 1/12 physics analysis and forecast updated daily product [data set]." [Online]. Available: https://doi.org/10.48670/moi-00016

[17] P. F. J. Lermusiaux *et al.*, "Quantifying uncertainties in ocean predictions," *Oceanography*, vol. 19, no. 1, pp. 92–105, 2006.

[18] P. F. J. Lermusiaux, "Uncertainty estimation and prediction for interdisciplinary ocean dynamics," *Journal of Computational Physics*, vol. 217, no. 1, pp. 176–199, 2006.

[19] N. Pinardi *et al.*, "The sea: The science of ocean prediction," *Journal of Marine Research*, vol. 75, no. 3, pp. 101–102, May 2017, special issue: The Science of Ocean Prediction, vol. 17 of The Sea.

[20] ——, "From weather to ocean predictions: an historical viewpoint," *Journal of Marine Research*, vol. 75, no. 3, pp. 103–159, May 2017, special issue: The Science of Ocean Prediction, vol. 17 of The Sea.

[21] T. Lolla *et al.*, "Path planning in time dependent flow fields using level set methods," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 166–173.

[22] D. N. Subramani *et al.*, "Time-optimal path planning: Real-time sea exercises," in *Oceans '17 MTS/IEEE Conference*, Aberdeen, Jun. 2017.

[23] P. F. J. Lermusiaux *et al.*, "Optimal planning and sampling predictions for autonomous and Lagrangian platforms and sensors in the northern Arabian Sea," *Oceanography*, vol. 30, no. 2, pp. 172–185, Jun. 2017, special issue on Autonomous and Lagrangian Platforms and Sensors (ALPS).

[24] W. Zhang *et al.*, "Optimal trajectory generation for a glider in time-varying 2d ocean flows b-spline model," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 1083–1088.

[25] M.-C. Tsou and H.-C. Cheng, "An ant colony algorithm for efficient ship routing," *Polish Maritime Research*, 2013.

[26] V. T. Huynh *et al.*, "Predictive motion planning for auvs subject to strong time-varying currents and forecasting uncertainties," in *International Conference on Robotics and Automation*, 2015, pp. 1144–1151.

[27] D. Kularatne *et al.*, "Going with the flow: a graph based approach to optimal path planning in general flows," *Autonomous Robots*, vol. 42, no. 7, pp. 1369–1387, 2018.

[28] G. Mannarini and L. Carelli, "VISIR-1. b: Ocean surface gravity waves and currents for energy-efficient navigation," *Geoscientific Model Development*, vol. 12, no. 8, pp. 3449–3480, 2019.

[29] D. Kularatne *et al.*, "Optimal path planning in time-varying flows with forecasting uncertainties," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 4857–4864.

[30] D. N. Subramani and P. F. Lermusiaux, "Risk-optimal path planning in stochastic dynamic environments," *Computer Methods in Applied Mechanics and Engineering*, vol. 353, pp. 391–415, 2019.

[31] A. A. Pereira *et al.*, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *Journal of Field Robotics*, vol. 30, no. 5, pp. 741–762, 2013.

[32] G. A. Hollinger *et al.*, "Learning uncertainty in ocean current predictions for safe and reliable navigation of underwater vehicles," *Journal of Field Robotics*, vol. 33, no. 1, pp. 47–66, 2016.

[33] G. Knizhnik *et al.*, "Flow-based control of marine robots in gyre-like environments," in *IEEE International Conference on Robotics and Automation*, 2022, pp. 3047–3053.

[34] P. Gunnarson *et al.*, "Learning efficient navigation in vortical flow fields," *Nature communications*, vol. 12, no. 1, pp. 1–7, 2021.

[35] W. H. Al-Sabban *et al.*, "Wind-energy based path planning for unmanned aerial vehicles using markov decision processes," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 784–789.

[36] A. Chakrabarty and J. Langelaan, "UAV flight path planning in time varying complex wind-fields," in *American Control Conference*, 2013, pp. 2568–2574.

[37] M. G. Bellemare *et al.*, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.

[38] D. Jones and G. A. Hollinger, "Planning energy-efficient trajectories in strong disturbances," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2080–2087, 2017.

[39] M. Otte *et al.*, "Any-time path-planning: Time-varying wind field+ moving obstacles," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 2575–2582.

[40] D. N. Subramani and P. F. J. Lermusiaux, "Risk-optimal path planning in stochastic dynamic environments," *Computer Methods in Applied Mechanics and Engineering*, vol. 353, pp. 391–415, Aug. 2019.

[41] N. K. Yilmaz *et al.*, "Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming," *IEEE Journal of Oceanic Engineering*, vol. 33, no. 4, pp. 522–537, Oct. 2008.

[42] ——, "Path planning methods for adaptive sampling of environmental and acoustical ocean fields," in *Oceans 2006*, Boston, MA, 18-21 Sept 2006, p. 6.

[43] A. Alexander *et al.*, "Loon library: Lessons from building loon's stratospheric communications service," Loon LLC, Tech. Rep., 2021. [Online]. Available: https://storage.googleapis.com/x-prod.appspot.com/files/The%20Loon%20Library.pdf

[44] B. Eysenbach *et al.*, "Search on the replay buffer: Bridging planning and reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[45] S. Pateria *et al.*, "Hierarchical reinforcement learning: A comprehensive survey," *ACM Comput. Surv.*, vol. 54, no. 5, jun 2021.

[46] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book, 2018.

[47] K. Doya, "Reinforcement learning in continuous time and space," *Neural computation*, vol. 12, no. 1, pp. 219–245, 2000.

[48] S. Bansal *et al.*, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *IEEE International Conference on Decision and Control*, 2017, pp. 2242–2253.

[49] I. M. Mitchell and J. A. Templeton, "A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems," in *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC*, 2005, pp. 480–494.

[50] J. Wu *et al.*, "Carbon dioxide removal via macroalgae open-ocean mariculture and sinking: An earth system modeling study," *Earth System Dynamics*, 2022.

[51] I. Martins and J. Marques, “A model for the growth of opportunistic macroalgae (enteromorpha sp.) in tidal estuaries,” *Estuarine, Coastal and Shelf Science*, vol. 55, no. 2, pp. 247–257, 2002.

[52] J. Zhang *et al.*, “A model for the growth of mariculture kelp saccharina japonica in sanggou bay, china,” *Aquaculture Environment Interactions*, vol. 8, 04 2016.

[53] E. P. Chassignet *et al.*, “US GODAE: global ocean prediction with the HYbrid Coordinate Ocean Model (HYCOM),” *Oceanography*, vol. 22, no. 2, pp. 64–75, 2009.

[54] M. Wiggert *et al.*, “Navigating underactuated agents by hitchhiking forecast flows,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, Cancún, Mexico, Dec. 2022, pp. 2417–2424.

[55] L. Beckenbach and S. Streif, “Approximate infinite-horizon predictive control,” in *IEEE International Conference on Decision and Control*, 2022, pp. 3711–3717.

[56] D. Silver *et al.*, “Mastering the game of go with deep neural networks and tree search,” *Nature*, vol. 529, pp. 484–503, 2016.

[57] D. N. Subramani *et al.*, “Stochastic time-optimal path-planning in uncertain, strong, and dynamic flows,” *Computer Methods in Applied Mechanics and Engineering*, vol. 333, pp. 218–237, 2018.

[58] P. F. J. Lermusiaux *et al.*, “Real-time probabilistic coupled ocean physics-acoustics forecasting and data assimilation for underwater GPS,” in *OCEANS 2020 IEEE/MTS*. IEEE, Oct. 2020, pp. 1–9.