

# RRT\*-Enhanced Long-Horizon Path Planning for AUV Adaptive Sampling using a Cost Valley

---

**Yaolin Ge**

Department of Mathematical Sciences  
Norwegian University of Science and Technology  
Trondheim, Norway  
yaolin.ge@ntnu.no

**Jo Eidsvik**

Department of Mathematical Sciences  
Norwegian University of Science and Technology  
Trondheim, Norway  
jo.eidsvik@ntnu.no

**André Julius Hovd Olaisen**

Department of Mathematical Sciences  
Norwegian University of Science and Technology  
Trondheim, Norway  
andre.j.h.olaisen@ntnu.no

## Abstract

With the goal of effective sampling of oceanographic variables by autonomous underwater vehicles, we propose a long-horizon adaptive sampling system which consists of a flexible cost valley concept and a non-myopic path planner. The goal of sampling is to gather information to reduce the expected variability or classification error of a particular feature. At the same time, one must adhere to operational challenges such as obstacles or time constraints. We combine various such objectives allocated to spatial locations in a weighted calculation to construct a cost surface, and the cost valley concept is then based on maneuvering to the minimal parts of this surface. We combine this cost valley concept with a rapidly exploring random trees strategy for non-myopic path planning. Results from a field trial in a Norwegian fjord show that the autonomous underwater vehicle running our long-horizon adaptive sampling system onboard successfully explores the salinity field in a river plume during a several-hours mission.

# 1 Introduction

Autonomous underwater vehicles (AUVs) have been used extensively to investigate different oceanographic phenomena. See e.g. Hwang et al. (2019) for a recent review. Benefits of AUV exploration are fast deployment and guided in-situ measurements without much human involvement. By bridging statistical modeling, embedded computing, and sensor technology, one can conduct AUV missions using adaptive sampling, which has gained interest in a variety of oceanographic application such as quantification of chlorophyll (Fossum et al., 2019; Zhang et al., 2020), detection of dissolved oxygen (Stankiewicz et al., 2021), benthic habitat mapping (Rigby et al., 2010; Anyosa et al., 2023) and frontal zones characterization (Zhang et al., 2016; Fossum et al., 2021; Ge et al., 2023).

In using only one AUV to conduct adaptive sampling, dominating methods can be grouped into either myopic (greedy) or non-myopic approaches. Myopic strategies guide the AUV towards the most informative location selected from a subset of candidate locations within the myopic (near-sighted) neighborhood radius (Fossum et al., 2021). The greediness of such computationally effective algorithms can make them fail at revealing new interesting areas that are not in the vicinity of the current location. Non-myopic strategies can alleviate such challenges by expanding their search horizon (Bai et al., 2021). Xiao and Wachs (2022) demonstrate long-horizon algorithms in a small-scale case study where the shapes of unknown objects are revealed by a robot arm, using prediction variance reduction as the main criterion. Suh et al. (2017) present an idea of using cost-aware rapidly exploring random trees (RRT\*) to generate sampling paths based on cross-entropy as a cost function. However, the computational costs associated with such non-myopic algorithms are usually very high. Although there are marine field applications using RRT\* (Enevoldsen and Galeazzi, 2021; Zacchini et al., 2022), it has shown difficult to apply methods such as RRT\* for large-scale onboard computations in the ocean. In particular, the tree expansion used in such algorithms is challenging in real-time operations, and with limited trees it loses some of its strengths. It is interesting to find ways to robustify the tree construction, for instance by combining multiple criteria in the objective function used for sampling.

Adaptive sampling relies on a criterion that allows the AUV to recognize high or low costs. In some sampling situations, one can explicitly relate the rewards or costs of sampling to an underlying decision situation, and then compute the value of information in monetary values (Eidsvik et al., 2015). In other situations, one can ask more expeditionary questions related to scientific task execution, such as Preston et al. (2022) who employ an AUV 'sniffing' for hydrothermal vents. But in many situations, the reward is more ambiguous and multifaceted, say via occupancy grids for presence or absence of a feature (Jakuba and Yoerger, 2008),

reduced variance or entropy of a field variable (Binney et al., 2013; Berget et al., 2018; Ma et al., 2018; Fossum et al., 2018) or by identifying hot spots in the field or the different water masses (Das et al., 2010; Fossum et al., 2021; Ge et al., 2023). Moreover, very high costs should be assigned to collisions or neglected time constraints. We will focus on a multi-objective cost function here, and in doing so we will incorporate both information rewards and operational costs. Jaillet et al. (2010) apply a cost map approach to optimize path planning in different case studies within a predefined configuration space. Similarly, Ettlín and Bleuler (2006) show the feasibility of utilizing a blended cost map based on the weighted sum of sub-components for rough terrain path planning using RRT. Okopal (2019) shows benefits of multiple objectives in the setting of evolving mission policies. The sub-component of the cost map can thus represent the terrain roughness or other operational aspects. However, the cost map’s inflexibility during operations limits its usefulness.

In this study, we design a long-horizon adaptive sampling system for AUVs for river plume front mapping. This system allows an AUV to adjust its paths based on samples, which is beneficial in mapping frontal zones. It is essential to assess the system’s ability to function in regions like river plume fronts, affected by ocean currents. Operational safety, budgetary concerns, collision hazards, and regional bathymetry are also considered. Given the river plume front’s dynamic nature, it is a suitable model to analyze system robustness in intricate settings. Our system comprises two main components: i) the cost valley, and ii) the path planner. The cost valley directs the AUV on the most efficient path for detecting the river plume front, incorporating collision avoidance within a set distance budget. This is achieved through weighted penalty and reward fields. For the path planning, we employ the RRT\* for long-horizon planning, and in doing so we demonstrate the algorithm’s efficiency in the Trondheim fjord deployment by refining an ocean model output using in-situ salinity measurements. The main contributions of this paper are:

- A versatile multi-criterion cost surface where its cost valley defines promising AUV sampling paths,
- A RRT\* algorithm for exploring the cost valley in a long-horizon planner,
- An AUV field deployment conducting adaptive sampling in a Norwegian fjord based on RRT\* and the cost valley concept including weighted information fields as well as collision and distance budget.

In Section 2 we describe the background which motivates the study of long-horizon path planning to map the river plume front. In Section 3 we outline our approach for onboard AUV computations. In Section 4 we present the cost valley concept. In Section 5 we show results from a simulation study. In Section 6 we show the results of the AUV field deployment. In Section 7 we conclude and point to future research directions.

## 2 Problem Statement

Frontal regions have been recognized as important for physical and biological oceanography. Fedorov (1986) discuss how fronts are critical for the dynamic and kinematic features arising when energy is transmitted in the oceanic basins. Belkin et al. (2009) describe frontal systems as hot-spots that shape parts of the marine ecosystem. We are particularly interested in river plumes in this paper. Rivers carry more than one-third of land-based precipitation to the ocean, and there has recently been much research on the transformation of river freshwater discharge and its dissolution into the ambient ocean, see e.g. Horner-Devine et al. (2015) for a review. This is important for instance in the context of agricultural run-off and understanding how riverine nutrients reach the open ocean (Sharples et al., 2017).

Figure 1 shows a satellite image illustrating typical water mass mixing during the Spring in the Trondheim fjord, Norway. The original image is enhanced by using the SNAP software developed by Zuhlke et al. (2015) which uses basic remote sensing techniques such as contrast stretching, color balancing, and false-color composites to highlight the visual features of water mass mixing between the ocean and the river (Richards and Richards, 2022). It is not obvious how to interpret the different colors in an image like Figure 1, even though one might be able to see river water masses with different sediment types in such a snapshot. To build more nuanced models, one typically also relies on numerical ocean models and in-situ measurements for calibration.

We zoom in on the Nidelva river (Figure 1) in this study, with the goal of characterizing the river plume region using ocean models and an AUV. Figure 2(a) shows the salinity field from numerical ocean model SINMOD (Slagstad and McClimans, 2005) on 11 May 2022. SINMOD <https://www.sintef.no/en/ocean/initiatives/sinmod/> is a multi-purpose numerical ocean model that can be connected to biological oceanography models with a broad spectrum of applications. For our situation, we focus only on the salinity output from SINMOD. The results in Figure 2 are constructed by averaging numerical ocean model data over time. It shows multiple river outlets causing regional river plumes.

For the AUV field deployment, the operational area is chosen based on the interesting parts of the outputs from the numerical ocean model, see Figure 2(b). The outer polygon draws the border for safe operation, whereas the inner polygon is an unsafe region due to shallow waters. For autonomous operation, we select a start location (small red dot to the east) and an end location (big blue dot), see Figure 2(b). The objective is then to conduct long-term AUV sampling without human intervention. Hence, the robot must travel from the start location to the end location to sample the salinity field informatively while returning home in

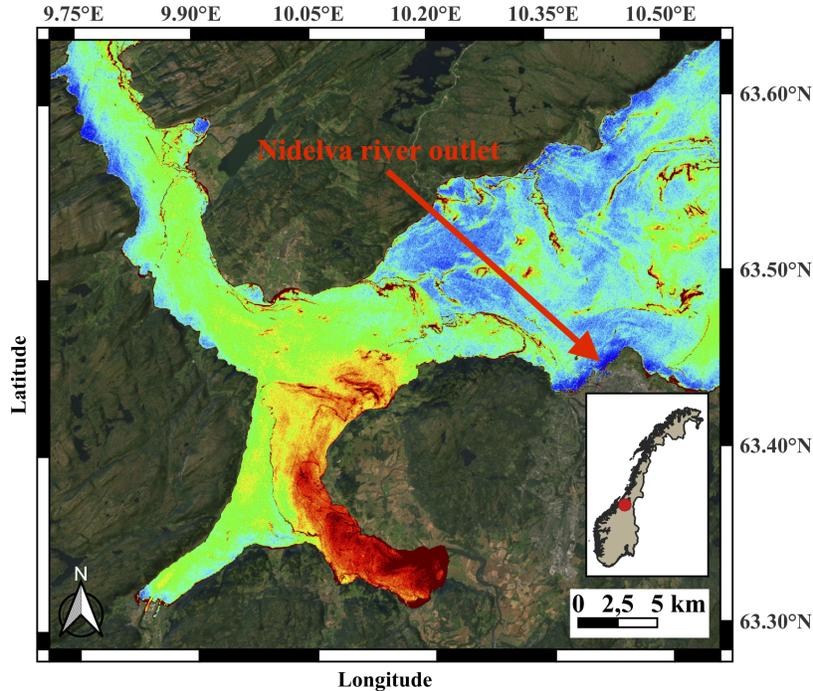


Figure 1: Sentinel-2 image of ocean water mixing on June-4th, 2021 in Trondheimsfjorden. Basic remote sensing techniques such as contrast stretching, color balancing, and false-color composites are applied. The arrow points out the Nidelva river outlet that we are interested in. Courtesy of the Copernicus, ESA

time and with operational constraints such as avoiding the shallow region near the island in the center and time or traveling distance constraints. In practice, the time constraints can be translated to the maximum allowance for the traveling range. Start locations and destination are chosen to enable reasonable mapping of the entire river plume frontal region, well within the specified distance budget.

### 3 Spatio-temporal model

We here define the Gaussian spatio-temporal surrogate model for the ocean salinity. The main benefit of having such a Gaussian surrogate model is that it enables real-time model updating onboard the AUV.

#### 3.1 Prior model

A spatio-temporal prior model for salinity is trained from SINMOD numerical ocean model data. This prior represents a statistical surrogate model of the complex physical processes in the ocean, that nevertheless mimics the key trends, variability and correlations of salinity in space and time for our case study. The prior model is here defined as a spatial auto-regressive Gaussian process model (Cressie and Wikle, 2015).

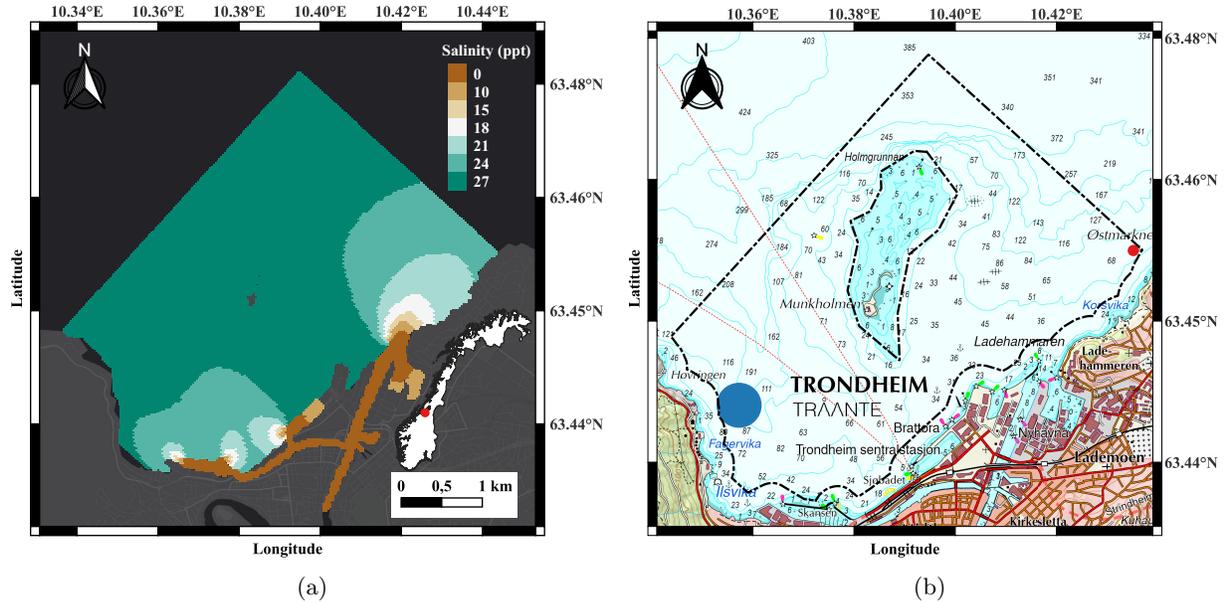


Figure 2: (a) Numerical ocean model data for surface salinity on May 11th, 2022. (b) The selected operational area, the outer polygon draws the border for the operation whereas the inner polygon shows an unsafe region. The small red dot shows where to start, and the big blue dot shows where to end.

We assume a spatially discretized domain of  $n$  locations;  $\mathbf{u}_1, \dots, \mathbf{u}_n$  where  $\{\mathbf{u} \in \mathcal{M} \subset \mathcal{R}^2\}$ . This grid covers a lateral domain, with depth fixed at 0.5m. Times are indicated by  $t = 0, 1, \dots$ , discretized in a regular sampling time interval. We denote the spatio-temporal salinity variable by  $\boldsymbol{\xi}_t = (\xi_{t,\mathbf{u}_1}, \dots, \xi_{t,\mathbf{u}_n})^T$ .

The initial state is Gaussian distributed  $\boldsymbol{\xi}_0 \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , where the length- $n$  mean vector  $\boldsymbol{\mu}$  is specified by averaging SINMOD numerical ocean model data over time at every location. For the  $n \times n$  covariance matrix  $\boldsymbol{\Sigma}$ , we assume constant variance  $\sigma^2$  and a Matérn correlation function so that covariance elements are  $\Sigma(i, i') = \sigma^2(1 + \phi h(i, i')) \exp(-\phi h(i, i'))$ , where the correlation decay parameter is  $\phi$  and with Euclidean distance  $h(i, i')$  between sites  $\mathbf{u}_i$  and  $\mathbf{u}_{i'}$ ,  $i, i' = 1, \dots, n$ . We specify parameters  $\sigma$  and  $\phi$  using variogram plots of SINMOD data (Cressie and Wikle, 2015).

The temporal variation is defined by an autoregressive process:

$$\boldsymbol{\xi}_t = \boldsymbol{\mu} + \rho(\boldsymbol{\xi}_{t-1} - \boldsymbol{\mu}) + \mathbf{v}_t, \quad \mathbf{v}_t \sim N(0, (1 - \rho^2)\boldsymbol{\Sigma}), \quad t = 1, \dots, \quad (1)$$

where the scalar autocorrelation parameter  $\rho$ . Assuming  $|\rho| \leq 1$ , this is a stationary process over time so that the marginal distribution at any time is  $\boldsymbol{\xi}_t \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  for  $t \geq 0$ . The extreme case with  $\rho = 1$  represents a spatial model without temporal variation. With  $\rho = 0$ , the spatial fields at different times  $t$  are uncorrelated. In our field study, the parameter  $\rho$  is trained from correlations over discretized time steps in the SINMOD

data for the same location.

### 3.2 Updating

The prior model described in equation (1) is updated by in-situ AUV measurements, where we then assume that the AUV is cruising in the lateral plane at 0.5 m depth. We model the AUV measurement  $y_t$  at stage or time  $t = 1, \dots$  by

$$y_t | \boldsymbol{\xi}_t \sim N(\mathbf{f}_t^T \boldsymbol{\xi}_t, r^2), \quad (2)$$

where the vector  $\mathbf{f}_t$  defines the spatial sampling indices at this stage of operation and  $r$  is the salinity measurement noise standard deviation.

The statistical surrogate model running onboard the AUV is updated with the in-situ salinity measurements. We apply Bayes' rule to achieve data assimilation at times  $t$ . Similar to a Kalman filter with the state vector now representing the spatial salinity field, this updating is done in real-time onboard the AUV. Between measurement times, the dynamical model propagates the state variable mean and covariance. Denoting the predictive mean and covariance by  $\mathbf{m}_{t|t-1}$  and  $\mathbf{S}_{t|t-1}$ , the updated mean  $\mathbf{m}_{t|t}$  and covariance  $\mathbf{S}_{t|t}$  are available by the recursive Kalman filter formulae defined by

$$\begin{aligned} \mathbf{m}_{t|t-1} &= \boldsymbol{\mu} + \rho(\mathbf{m}_{t-1|t-1} - \boldsymbol{\mu}) \\ \mathbf{S}_{t|t-1} &= \rho^2 \mathbf{S}_{t-1|t-1} + (1 - \rho^2) \boldsymbol{\Sigma} \\ \mathbf{G}_t &= \mathbf{S}_{t|t-1} \mathbf{f}_t (\mathbf{f}_t^T \mathbf{S}_{t|t-1} \mathbf{f}_t + r^2)^{-1} \\ \mathbf{m}_{t|t} &= \mathbf{m}_{t|t-1} + \mathbf{G}_t (y_t - \mathbf{f}_t^T \mathbf{m}_{t|t-1}) \\ \mathbf{S}_{t|t} &= \mathbf{S}_{t|t-1} - \mathbf{G}_t \mathbf{f}_t^T \mathbf{S}_{t|t-1}. \end{aligned} \quad (3)$$

Here, we start by  $\mathbf{m}_{1|0} = \boldsymbol{\mu}$  and  $\mathbf{S}_{1|0} = \boldsymbol{\Sigma}$  at the first step.

If the AUV is pausing, the last three steps in equation (3) do not take place, as there is no data updating. In that situation, one will just propagate the mean and covariance expressions according to the first two steps.

Regarding scalability, the Gaussian updating formula in equation (3) requires matrix factorization (inverse calculation) of a matrix with size equal to the amount of data gathered at each time point. In our case with sequential data assimilation, only data  $y_t \in \mathcal{R}$  at a single waypoint node is included, and this factorization is hence very fast. A bigger challenge here is the evaluation and storage of the  $n \times n$  covariance matrices  $\mathbf{S}_{t|t-1}$  and  $\mathbf{S}_{t|t}$ , but for the waypoint graph in the two-dimensional longitude-latitude domain, this does not

cause challenges.

## 4 Long-horizon path planning using cost valley

In this section, we introduce our concept for constructing a cost valley in the context of long-horizon path planning. A cost valley refers to a cost function that takes the shape of a valley when plotted on a 3D graph (see Figure 5). This shape results from certain locations having lower costs than others, thus creating a valley-like structure in the cost function.

Identifying the region of the cost function that represents the most efficient path is crucial for designing algorithms that can help robots navigate through complex environments in the most efficient way possible. To this end, the agent can utilize different planners such as the myopic planner (representing a greedy approach) or the RRT\* planner (representing a long-horizon search). Each planner has its own advantages and disadvantages, so it is up to the designer to choose depending on the specific application.

Our approach for constructing a cost valley involves two groups of cost sub-fields: i) operational and ii) informative. The operational cost fields guide the agent safely and efficiently to its destination, while the informative cost fields help the agent make informed decisions about where to sample by highlighting areas of information gain returning a reward. We explain each group of cost fields in the following sections to better understand how they are used to construct the cost valley.

Once the cost valley is defined, we put this in a path planning framework and explain our system architecture for using this in a field robotics setting.

### 4.1 Operational cost fields

The operational cost fields are designed to guide the agent responsibly while ensuring that it reaches its destination on time. This is achieved through the use of two different cost sub-fields: obstacle avoidance and budget cost fields. The obstacle avoidance cost field prevents the agent from colliding with obstacles. In our case, it gives infinite cost for the island region in the Trondheim fjord and zero cost elsewhere. The time budget cost field ensures that the AUV reaches its destination before the time or distance budget runs out. It assigns a large cost to areas outside of the budget ellipse and zero cost to areas inside it. Both these cost fields work together to constrain the agent within a specific operational frame, as shown in Figure 3.

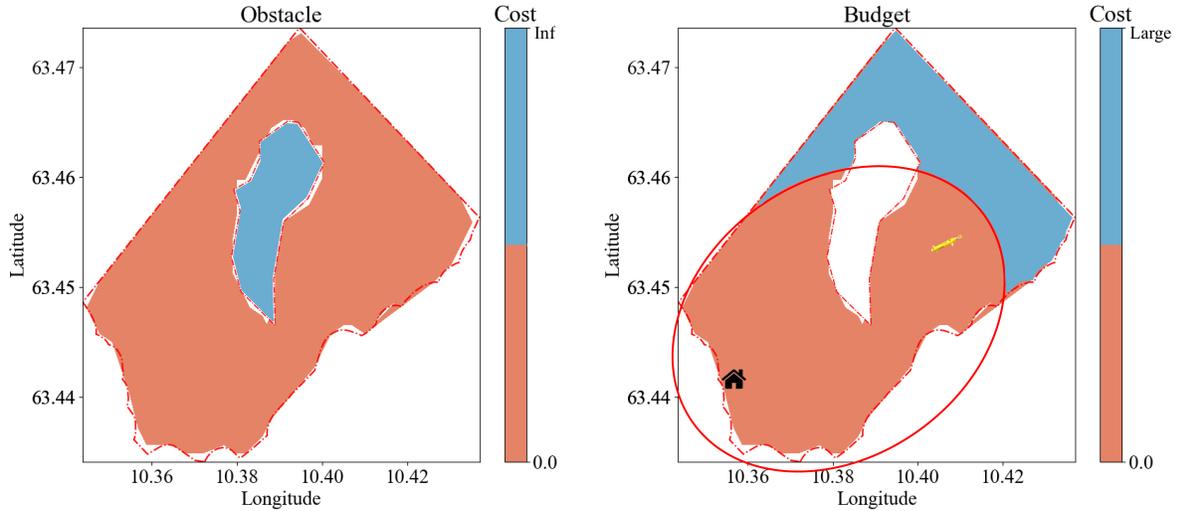


Figure 3: Operational cost fields. The left display shows an obstacle cost field that assigns infinite cost to the island in the middle of the operation domain and zero cost to other areas. The right display shows a budget cost field that assigns high costs to areas outside the budget ellipse and zero costs to areas inside. Here, the AUV should go to the end location (house) before the distance budget runs out.

## 4.2 Informative cost fields

The informative cost fields in our approach are focused on aiding in efficient sampling of the river plume front. We use two criteria; integrated variance reduction (IVR) and expected integrated Bernoulli variance (EIBV), which are incorporated into the cost field, as shown in Figure 4.

The IVR cost field aims to reduce the overall uncertainty of the field. For a given sampling design, the IVR is the trace (sum of diagonal elements) of the covariance reduction matrix  $\mathbf{R} = \mathbf{G}_t \mathbf{f}_t^T \mathbf{S}_{t|t-1}$  in equation (3). Here, the design enters in the vector  $\mathbf{f}$ . Note that the cost related to IVR is inversely proportional, meaning that larger IVR results in smaller costs, and vice versa. Sampling locations that have been visited recently, will give a small IVR. Sampling locations that are yet unexplored or have not been explored in a long time will tend to give large IVR. To reduce this IVR cost field, one would naturally sample the latter kinds of design for this criterion.

The EIBV cost field aims to guide the agent to locations where the river plume front may appear. The EIBV is in our case based on excursion probabilities related to a threshold and the associated Bernoulli variation. We let  $\zeta$  be a threshold in salinity. This threshold is chosen as it separates fresh river water from the more saline fjord water. At a location,  $\mathbf{u}_i$  and at time  $t$ , the excursion probability that salinity exceeds the threshold is then  $p(\mathbf{u}_i) = P(\xi_{t,\mathbf{u}_i} < \zeta) = \Phi_1 \left( \frac{\zeta - m_{t|t}(i)}{\sqrt{S_{t|t}(i,i)}} \right)$ , where  $\Phi_1$  denotes the univariate cumulative distribution function (CDF) of the standard normal. The Bernoulli variance is  $p(\mathbf{u}_i)(1 - p(\mathbf{u}_i))$ . With a goal

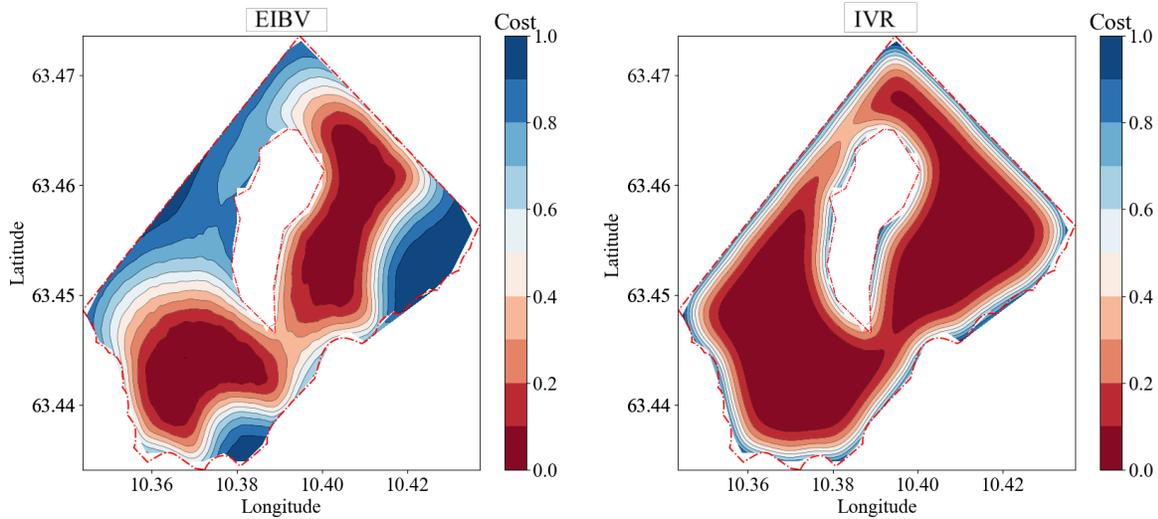


Figure 4: Informative cost fields at the initial stage. The left display shows the EIBV cost field, which assigns low cost to areas where the river plume front might exist at this initial time. The right display shows the IVR cost field, which assigns low cost almost everywhere as it has not yet sampled any in-situ measurements.

of classifying water masses, one would collect AUV data at sampling locations that reduce this Bernoulli variance. Unlike the IVR, the reduction will now depend on the observed data, but the expected reduction has a closed-form solution via a bivariate CDF (Ge et al., 2023). The cost field of EIBV of course depends on previously sampled locations. It further tends to resemble the spatial distribution of the river plume boundary given by the certain threshold  $\zeta$  because locations close to the threshold have a large Bernoulli variance, and therefore carry much information, see Figure 4.

Algorithm 1 shows how we compute these informative cost fields. Note how the design  $\mathbf{f}$  varies in the loop over spatial sampling locations  $i$ . The cost field in this way indicates the information value of the sampling individual locations. Each spatial site  $\mathbf{u}_i$  has an associated cost value for each criterion. The closed form solutions for both IVR and EIBV provided in Algorithm 1 ensure that the computations are relatively fast to do on the onboard computing units.

### 4.3 Cost valley construction

To compute the cost valley, we construct a weighted sum of all components using

$$\mathbf{C} = \mathbf{C}_{\text{obstacle}} + \mathbf{C}_{\text{budget}} + w_1 \cdot \mathbf{C}_{\text{EIBV}} + w_2 \cdot \mathbf{C}_{\text{IVR}}, \quad w_1 + w_2 = 1. \quad (4)$$

---

**Algorithm 1** Calculate informative cost fields for EIBV and IVR

---

**Require:**  $m_{t|t-1}, \mathbf{S}_{t|t-1}$   
 $\mathbf{EIBV} = \mathbf{0}^{n \times 1}$   
 $\mathbf{IVR} = \mathbf{0}^{n \times 1}$   
**for**  $i \in 1 \dots n$  **do**  
 $\mathbf{f} = \mathbf{0}^{n \times 1}, \mathbf{f}(i) = 1$   
 $\mathbf{R} = \mathbf{S}_{t|t-1} \mathbf{f} (\mathbf{f}^T \mathbf{S}_{t|t-1} \mathbf{f} + r^2)^{-1} \mathbf{f}^T \mathbf{S}_{t|t-1}$   
 $\mathbf{IVR}(i) = \sum_{i'=1}^n \mathbf{diag}(\mathbf{R})$   
 $\mathbf{EIBV}(i) = \sum_{i'=1}^n \Phi_2 \left( \begin{bmatrix} \zeta \\ -\zeta \end{bmatrix}; \begin{bmatrix} m_{t|t-1}(i') \\ -m_{t|t-1}(i') \end{bmatrix}, \mathbf{W}(i', i') \right),$   
where,  $\mathbf{W}(i', i') = \begin{bmatrix} S_{t|t-1}(i', i') & -R(i', i') \\ -R(i', i') & S_{t|t-1}(i', i') \end{bmatrix}$   
**end for**  
 $\mathbf{C}_{\text{EIBV}} = (\mathbf{EIBV} - \min(\mathbf{EIBV})) / (\max(\mathbf{EIBV}) - \min(\mathbf{EIBV}))$   
 $\mathbf{C}_{\text{IVR}} = \mathbf{1} - (\mathbf{IVR} - \min(\mathbf{IVR})) / (\max(\mathbf{IVR}) - \min(\mathbf{IVR}))$   
**return**  $\mathbf{C}_{\text{EIBV}}, \mathbf{C}_{\text{IVR}}$

---

The operational cost fields, which include the obstacle cost field  $\mathbf{C}_{\text{obstacle}}$  and the budget cost field  $\mathbf{C}_{\text{budget}}$ , are of utmost importance due to their ability to protect operational integrity. This is done by imposing hefty penalties in areas that are dangerous or far to reach. As a result, the EIBV and IVR cost fields ( $\mathbf{C}_{\text{EIBV}}$  and  $\mathbf{C}_{\text{IVR}}$ ), which are informative cost fields, have little effect on guiding the agent when compared to the strong penalties of the operational cost fields. Nevertheless, when operational safety is guaranteed and there is plenty of time, the costs in legitimate regions are minimal. In such cases, the main costs come from the informative cost fields, which direct the agent based on knowledge. In our approach, no weights are assigned to the operational cost fields to guarantee the agent's prompt reaction when any of these fields are activated. For example, it is essential that the agent returns to the base when the budget cost field is nearly exhausted. Therefore, locations located outside the allocated budget ellipse are heavily penalized, regardless of their perceived worth in the informative cost fields.

Figure 5 illustrates an instance of the cost valley, where the budget remains ample and the weights are evenly distributed ( $w_1 = w_2 = 0.5$ ). At each stage  $t = 1, \dots, N_{\text{steps}}$ , the AUV calculates this kind of cost valley by evaluating all criteria for all locations and weighting as in equation (4). The weighting aids the AUV sampling in balancing exploration of uncertain locations and clearly detecting the salinity boundaries, while maintaining the operational constraints for the vehicle.

#### 4.4 Path planning using the cost valley

We design a path planning algorithm to guide the agent in the adaptive sampling process. For this purpose we use the cost fields, and at each iteration, the least-cost path in the cost valley is calculated. This leads

the AUV from its current location to the one with the lowest cost. In doing so, one finds the best design  $\mathbf{D}_t$  at this stage in time  $t$ .

The algorithm we use for long-horizon path planning is described in Algorithm 2. To determine the optimal path, we here utilize the RRT\* path planner, as described in Karaman and Frazzoli (2011) and Hollinger and Sukhatme (2014). This planner computes the least-cost trajectory  $\mathbf{T}_t$  from the current location to the target location which is therefore the global minimum cost location, and the first location on this trajectory is selected as the next optimal design location  $\mathbf{D}_t$ . The AUV then takes in-situ measurements  $y_t$  at this location, and the model is updated using equation (3). This, in turn, updates the cost valley. The process continues until the budget is exhausted, with each new starting location being provided to the RRT\* path planner. In our approach, we also introduce a nuanced modification by adding an additional waypoint, termed the 'pioneer waypoint'. Initially, the AUV calculates two waypoints. After transitioning from the current waypoint to the next waypoint, the AUV immediately proceeds from this next waypoint to the pioneer waypoint. Concurrently, it performs computations for the forthcoming waypoints originating from the pioneer waypoint. In doing so, it enables us to do the real-time path planning. For a detailed explanation of using RRT\* path planner to determine the next waypoint, please refer to Ge et al. (2022).

In Figure 5, a 3D visualization is presented for one-step planning. Here the RRT\* planner uses the equal-weighted cost valley as a guide. As depicted in the display, the planned path (blue) that is computed from this cost valley facilitate long-horizon planning, and the planner is able to navigate away from the high-cost area. However, due to insufficient convergence in the RRT\* path planner, the path taken is sub-optimal, as evidenced by the detour in the trajectory. Increasing the number of iterations in the tree expansion can improve the optimality of the path. However, one must consider the trade-off between the computational cost and path optimality.

## 4.5 System architecture

Figure 6 shows how the architecture combines RRT\* path planning with cost valley guidance for adaptive sampling. The diagram also indicates how SINMOD is used to create an initial prior for the system, which is then improved using pre-survey data. The GRF (Gaussian random field) module is used to incorporate in-situ data into and updated state model including an uncertainty map, which is essential for the cost valley calculation. After this model update, an optimal trajectory is determined from the current location to the minimum cost location, producing the next waypoint. The ROS-IMC bridge is then used to send instructions to the AUV, which is discussed in more detail in the field deployment description. DUNE executes the control

## RRTStar and Cost valley illustration

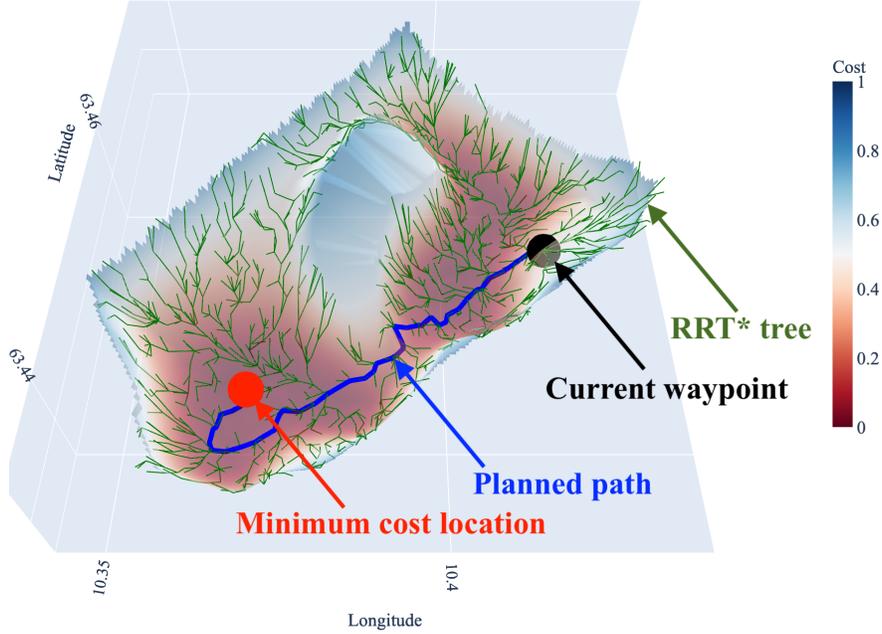


Figure 5: RRT\* path planning using cost valley illustrated in 3D. The green trees are generated using RRT\*, and the blue path is selected based on the minimum cost criterion. It leads the agent toward the global minimum cost location, shown as the red dot. The current waypoint is depicted as a black dot. One can see that RRT\* is not perfect and has a detour in the low-cost area close to the target location.

---

### Algorithm 2 Informative long-horizon path planning algorithm

---

**Require:** Initial mean  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ . Set  $\mathbf{m}_{1|0} = \boldsymbol{\mu}$  and  $\mathbf{S}_{1|0} = \boldsymbol{\Sigma}$ .

Set start waypoint  $\mathbf{D}_0$

$\mathcal{Y}_0 = \emptyset$

Budget = MaxDistanceRange

$t = 1$ .

**while** Budget  $\geq$  allowance **do**

**Plan :**

$\mathbf{CV} = \text{updateCostValley}(\mathbf{m}_{t|t-1}, \mathbf{S}_{t|t-1}, \text{Budget}, \mathbf{D}_{t-1})$

$\mathbf{u}_t = \text{argmin}_{\mathbf{u} \in \mathcal{M}}(\mathbf{CV})$

$\mathbf{T}_t = \text{RRT}^*(\mathbf{CV}, \mathbf{D}_{t-1}, \mathbf{u}_t)$

$\mathbf{D}_t = \mathbf{T}_t\{1\}$

    Budget = Budget -  $\|\mathbf{D}_t - \mathbf{D}_{t-1}\|_2$

**Act :**

    Go to waypoint  $\mathbf{D}_t$ .

**Sense :**

    Gather data  $y_t$ .  $\mathcal{Y}_t = (\mathcal{Y}_{t-1}, y_t)$ .

$\mathbf{G}_t = \mathbf{S}_{t|t-1} \mathbf{f}_t (\mathbf{f}_t^T \mathbf{S}_{t|t-1} \mathbf{f}_t + r^2)^{-1}$ .

$\mathbf{m}_{t|t} = \mathbf{m}_{t|t-1} + \mathbf{G}_t (y_t - \mathbf{f}_t^T \mathbf{m}_{t|t-1})$ .

$\mathbf{S}_{t|t} = \mathbf{S}_{t|t-1} - \mathbf{G}_t \mathbf{f}_t^T \mathbf{S}_{t|t-1}$ .

$\mathbf{m}_{t+1|t} = \boldsymbol{\mu} + \rho (\mathbf{m}_{t|t} - \boldsymbol{\mu})$ .

$\mathbf{S}_{t+1|t} = \rho^2 \mathbf{S}_{t|t} + (1 - \rho^2) \boldsymbol{\Sigma}$ .

$t = t + 1$ .

**end while**

---

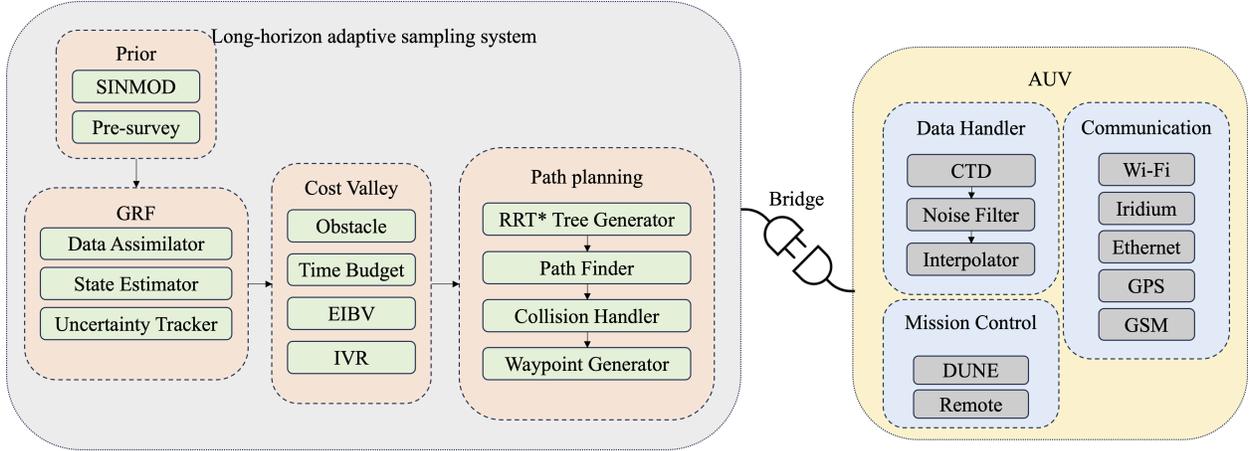


Figure 6: System architecture of our proposed adaptive sampling system using RRT\* for path planning and a cost valley for guidance.

command, allowing the AUV to sample autonomously.

## 5 Simulation study

In this simulation study, we explore the influence of various weighting schemes on AUV pathways. We also compare the RRT\* planner with its myopic counterpart. In doing so, we assess the system’s long-term robustness in identifying spatio-temporal complexities in a river plume system.

### 5.1 Simulation setup

Figure 3 shows the operational cost fields. They should always be given the highest priority to ensure adherence to safety constraints. Thus, the planner must evaluate the operational cost fields before evaluating the informative cost fields. The weights used should not affect the priority given to the operational cost fields. Therefore, we choose to study the planner’s behavior under different weighting schemes solely with respect to the informative EIBV and IVR cost fields (see Figure 4).

We can assign various weights between EIBV and IVR cost fields. It is important to note that the appropriate weight may vary depending on the application, and therefore, we cannot provide an ideal weighting scheme for all scenarios. In our study, we aim to demonstrate how extreme weighting schemes impact the system’s performance. To achieve this, we have created three cost scenarios with the following weights:

- EIBV dominant:  $w_1 = 1, w_2 = 0$

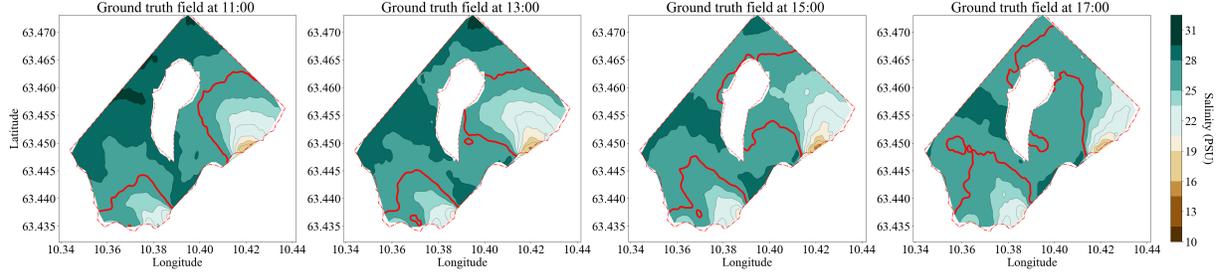


Figure 7: The surface salinity ground truth field at different time steps. The red contour shows the boundary of different water masses classified by the threshold.

- IVR dominant:  $w_1 = 0, w_2 = 1$
- Equal weight:  $w_1 = 0.5, w_2 = 0.5$

Replicate data sets are simulated using a temporal benchmark field of 8 hours from SINMOD as the mean value field. For the variation in the replicate runs, the Matérn covariance matrix coefficients are specified as  $\sigma = 1.0$  and  $\phi = 0.0064$ . The average salinity over the 100 replicates is shown in Figure 7 for different time steps. The red contour line delineates the boundary between saline and fresh water masses and is determined by a salinity threshold. It is clear that the numerical solution captures the tidal cycle as the boundary expands with the tide going from high to low, meaning the boundary is shifting outward as the current brings more freshwater from the river mouth.

Using the weight sets mentioned above, we construct three cost scenarios or valleys. For each cost valley, we run two agents starting at coordinates  $N63.440, E10.356$  in the WGS84 coordinate system. Two agents are employed, one using a myopic planner and the other a long-horizon RRT\* planner. Both have access to the cost valley field; however, the myopic agent is only able to assess the cost of adjacent locations from the cost valley and then select the next waypoint with the lowest cost. Thus, it does not take into account long-term considerations. Both agents have the same step size of 240m and are allowed to run for 8 hours, which is equivalent to 120 steps. At each sampling point, the AUV data is extracted from the replicate field data at the present location. We set measurement noise standard deviation  $r = 0.5$ .

## 5.2 Agent analysis for EIBV dominant field

We start by assessing typical AUV paths for a single replicate only. For clarity, we focus on the EIBV dominant cost valley scenario, contrasting the trajectories of the myopic and RRT\* planners under identical ground truth.

### 5.2.1 Myopic

Figure 8 shows the mean salinity field (left), uncertainty field (center), and cost valley (right) over time, following adaptive sampling with the myopic strategy across time iterations. Figure 8a showcases the fields after limited exploration near the start to the south-west. The cost valley directs the agent towards boundary regions, with the uncertainty field illustrating decreased uncertainty in visited areas. The eastern parts of the cost valley have low costs, but it is separated from the current AUV position by a high-cost region.

Figures 8b and 8c depict later stages with the agent predominantly exploring the field in the south-west. As time progresses, the boundary expands outward with the tide, and the agent’s sampling interest shifts in this direction. By 16:40, guided by the cost valley, the agent gravitates eastward. However, the low-cost region is now quite far to the north-east because of the tide, and clearly, time influences the agent’s posterior belief. Towards the end of the operation (Figure 8d), the east remains largely unexplored, and the boundary growth over time underscores the spatio-temporal model’s advantage in accommodating field dynamics. The myopic approach, with its limited foresight, often results in localized planning. Such agents risk entrapment in proximate regions, underscoring the need for expansive strategies. Nevertheless, the AUV has mapped the south-east parts of the river plume front very accurately.

### 5.2.2 Long-horizon RRT\*

Figure 9 presents the evolving mean salinity field, uncertainty field and cost valley from left to right for the long-horizon RRT\* strategy. The cost valley plots also display RRT\* trees, elucidating the planning process. The agent efficiently targets high-value regions for sampling.

Figure 9a captures the early phase of sampling, where the cost valley reveals the agent’s awareness of a prospective low-cost region in the east and its intention to navigate towards it. At the intermediate stages (Figures 9b and 9c), the AUV now moves between west and east parts of the plume front. Hence, in contrast to the myopic approach in Figure 8, the adaptive agent utilizing the long-horizon RRT\* strategy exhibits a sampling behavior that is covering much more of the spatial domain. Running RRT\* onboard, the agent’s ability to escape local attractions is clearly increased, and it spans longer to find regions offering maximum reward. Similar to what was seen for the myopic case, the AUV tends to move further out with the tide. In fact, in Figure 9d, in the concluding phase, the AUV judiciously navigates from the north past the obstacle to access the broader northern boundary. Overall, the RRT\* planner navigates the agent through dynamic terrains, and it spans much wider than the myopic planner. Even so, it could of course perform worse than

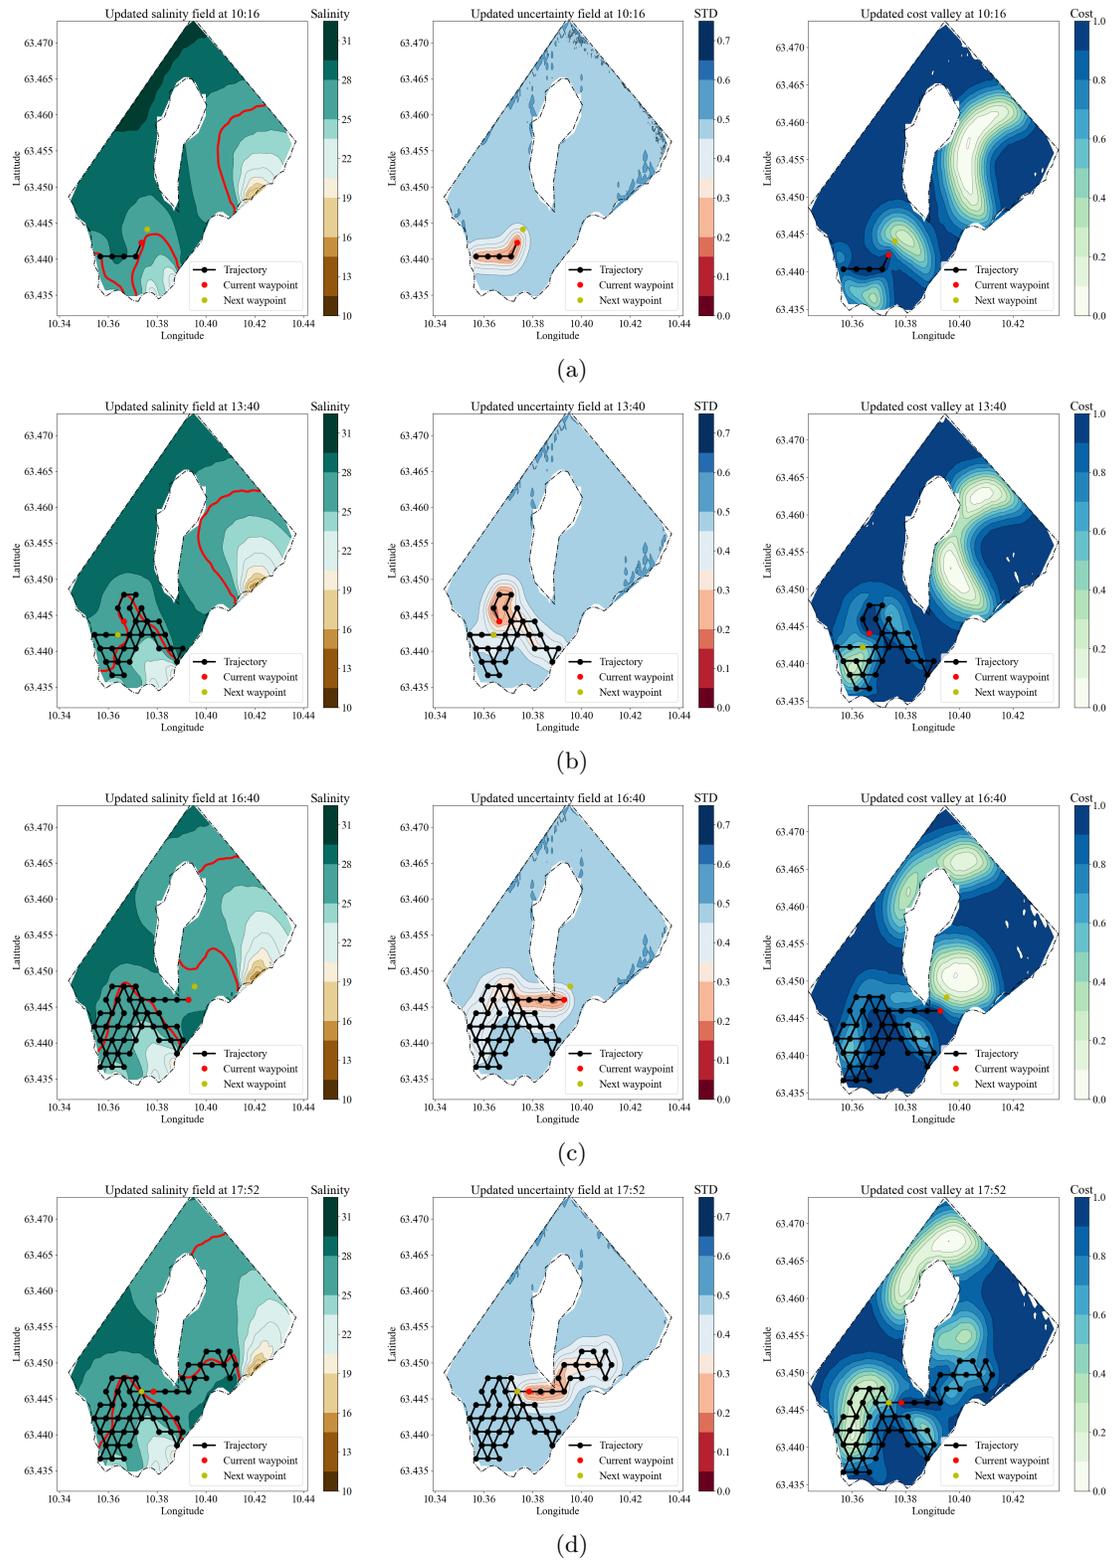


Figure 8: The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using the myopic strategy over a series of time iterations. In this particular example, most of the effort was focused on the southwest side before the agent moved eastward.

the myopic planner in detecting local plume details.

### 5.3 Traffic density analysis

To visualize the AUV sampling effort distribution across the 100 replicates, we devise a traffic flow density plot using kernel density estimation. This plot quantifies the intensity of AUV trajectories in various locations, pinpointing high-traffic areas. To underscore the influence of the cost valley, we employ this method for every weight set and for both the myopic and RRT\* planners.

#### 5.3.1 EIBV dominant cost valley

The EIBV cost valley prioritizes areas likely harboring the river plume front, as illustrated in Figure 4 (left). Often, these fronts manifest near river outlets. Figures 10a and 10b show the traffic-flow density maps for the myopic and RRT\* planners, respectively. At 11:52, the RRT\* strategy directs spread the AUV efforts toward key boundary regions, whereas the myopic planner’s efforts are limited to proximate areas around the starting location. It is not until 15:52 that the myopic planner starts to recognize the valuable eastern region in some of the replicate runs, a zone that the RRT\* planner has been covering extensively for hours. By 17:52, the myopic planner still has not sufficiently sampled the eastern region. For instance, it is missing the plume front at this time which is further out north because of the tide. The RRT\* planner has gained temporal insights in the temporal variations in this eastern part of the river plume front.

In practical terms, the RRT\* planner’s intentional traffic flow optimally directs the agent to efficiently sample key EIBV zones. For gathering equivalent data, it is preferable to expend minimal effort. The myopic planner, however, might become ensnared in particular local hot-spots. This underscores the importance of a well-conceived cost valley design to prevent over-exploitation of specific areas, ensuring a balanced sampling strategy.

#### 5.3.2 IVR dominant cost valley

For the IVR dominant cost valley, the objective is to maximize uncertainty reduction across the field. The traffic-flow density maps for both the myopic and RRT\* planners are portrayed in Figures 11a and 11b respectively.

Interestingly, while both planners initially adopt distinct paths, they soon exhibit comparable behaviors,

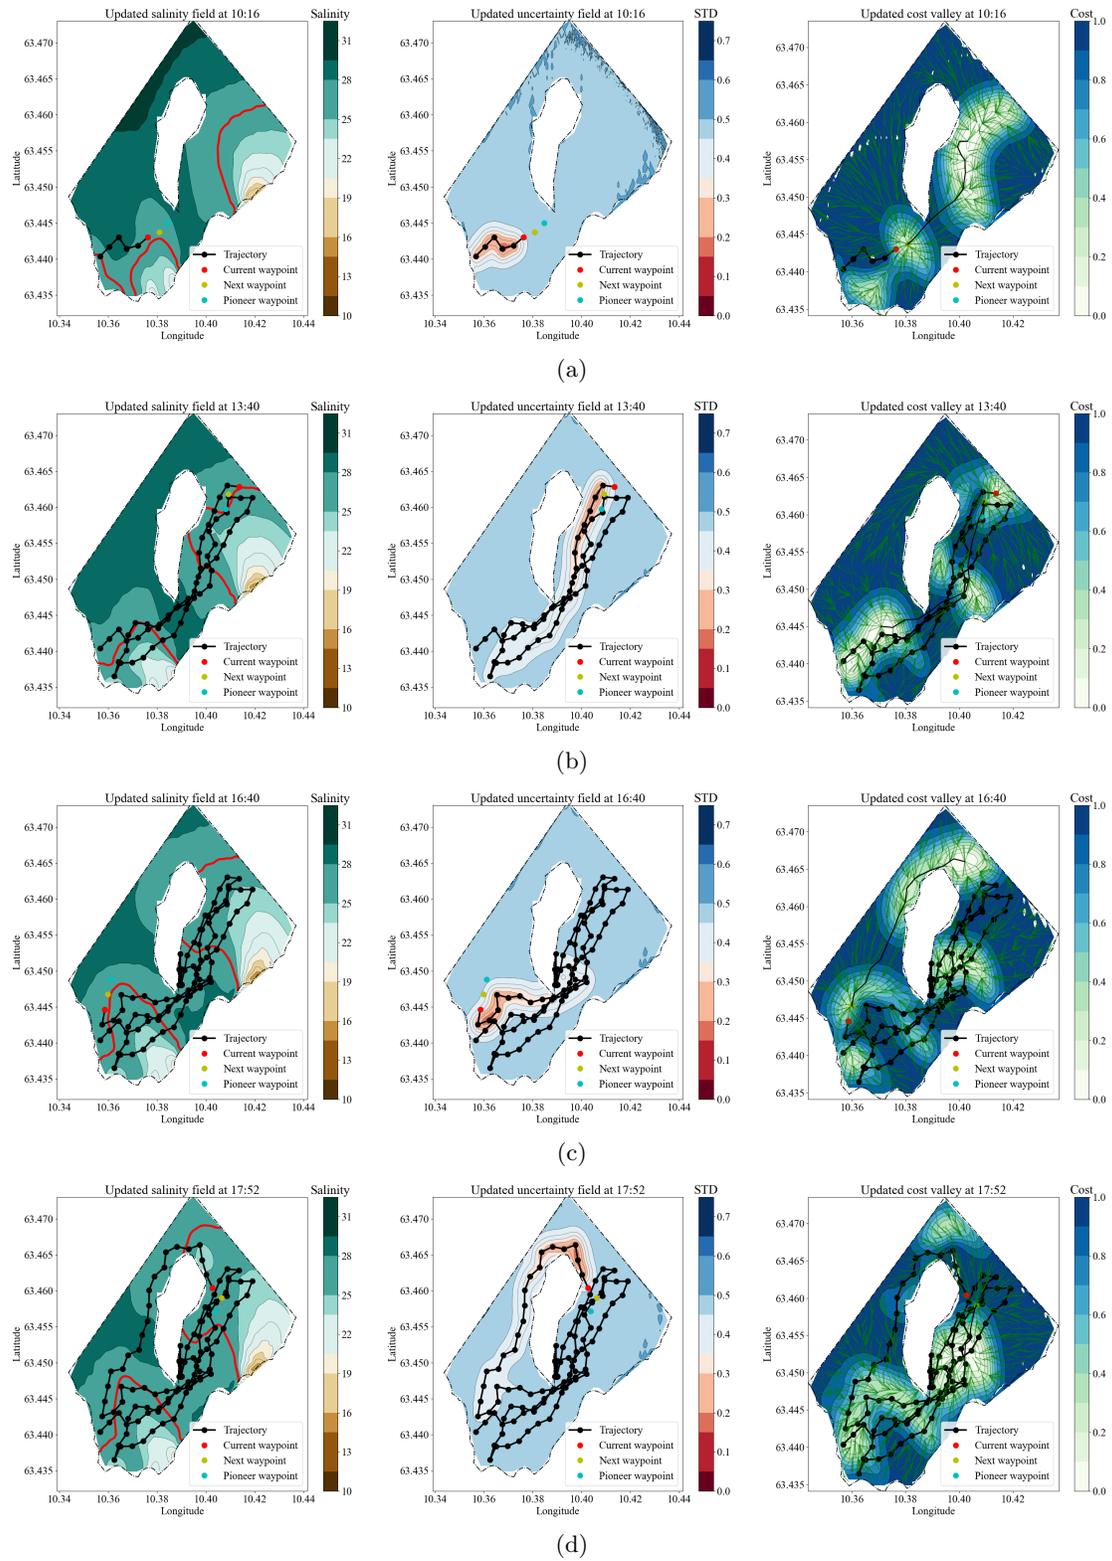


Figure 9: The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using our proposed long-horizon RRT\* strategy over a series of time iterations. The RRT\* trees are displayed in the cost valley column as well to better illustrate the planning mechanisms.

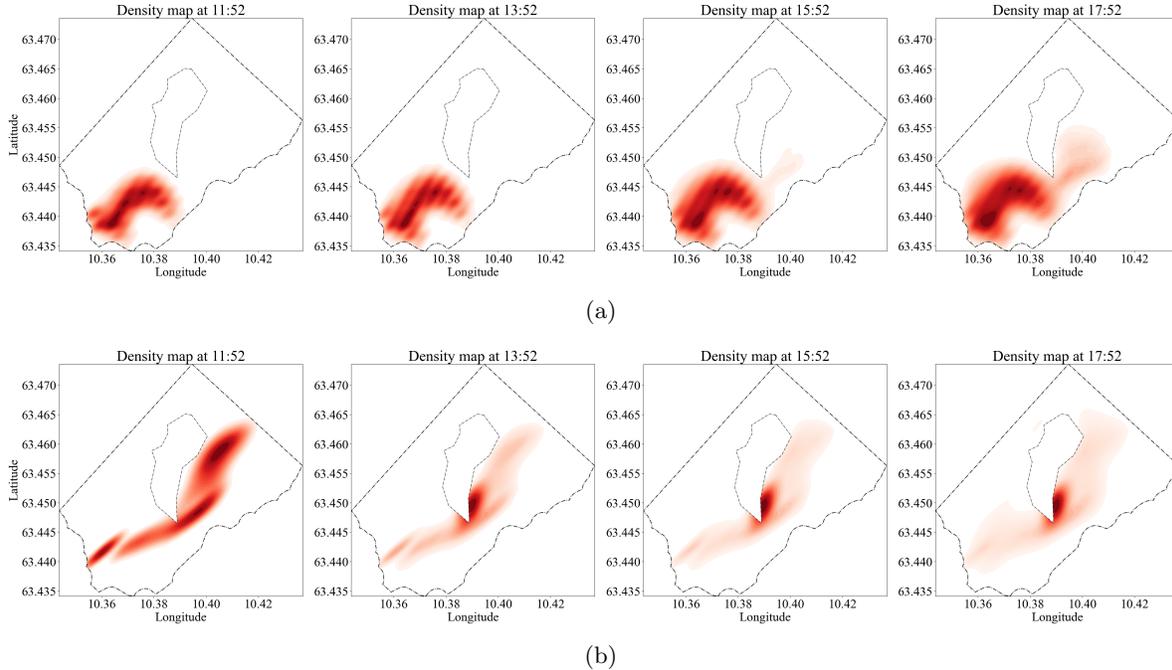


Figure 10: The traffic flow density map for the myopic agent (a) and the long-horizon RRT\* agent (b) using the EIBV dominant cost valley.

rapidly dispersing across the field. Consequently, distinguishing between the two planners becomes challenging in this context. They both effectively and expediently spread to cover the field comprehensively. Contrary to the EIBV cost valley, the IVR cost field depends solely on sampling locations in this situation with a Gaussian model, and it is not influenced by the data gathered at this location.

### 5.3.3 Equal weight cost valley

In this scenario, striking a balance between exploration (seeking unknown regions) and exploitation (sampling river plume fronts) is crucial. The intention is for the AUV to prioritize areas with low costs for both variance and the river plume fronts. The side-by-side traffic-flow density maps provided in Figures 12a and 12b offer a telling comparison between the myopic and RRT\* planners.

The myopic planner, particularly during its early hours of operation, displays a tendency to focus its sampling around its initial starting point, but less so than for the EIBV dominant strategy in Figure 10. It takes a longer time for the myopic strategy to expand its sampling reach towards the east. Its movements around an identified obstacle are quite balanced: it is about as likely to take a route south of the obstacle as it is to go north.

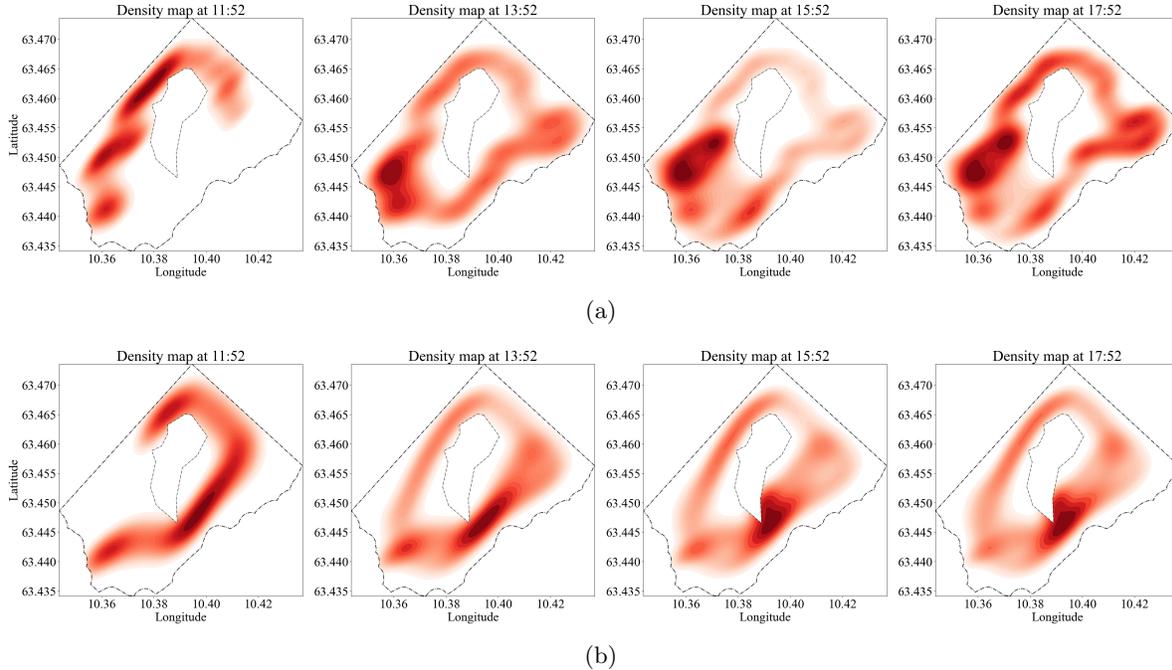


Figure 11: The traffic flow density map for the myopic agent (a) and the long-horizon RRT\* agent (b) using the IBV dominant cost valley.

Conversely, the RRT\* planner exhibits a more strategic approach that reflects the EIBV traffic maps in Figure 10 but more spread out to also cover the domain and reduce uncertainty. It is predominantly focusing on the southern regions for the majority of its operation. However, in the later hours, especially as the tidal effects become more prominent, the RRT\* planner directs its efforts towards the northern areas, albeit sparingly. This particular behavior underscores the planner’s adaptability to dynamic environmental changes such as tides.

When observing the coloration of the traffic-flow density maps, the myopic planner’s sampling pattern is evident. The darker red hues, particularly near its starting point, highlight its increased sampling frequency in that locale. This implies that the myopic strategy might be less efficient in adapting to changing environments or leveraging information from larger field areas. Not so for the RRT\* planner which has its more evenly spread hue, illustrating a more balanced and extensive coverage of the field.

#### 5.4 Metric analysis

In the replicate simulations, we monitor four key metrics: integrated Bernoulli variance (IBV), variance reduction (VR), root mean squared error (RMSE) and classification error (CE). For each metric, we take the average over all spatial grid cells, at every time point.

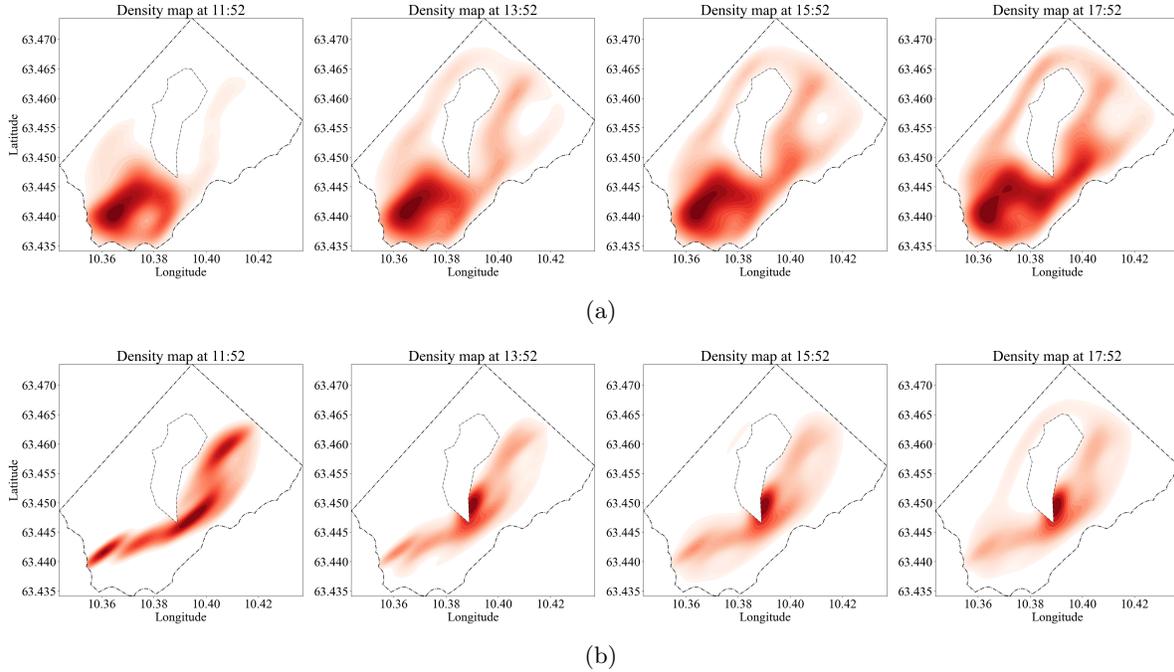
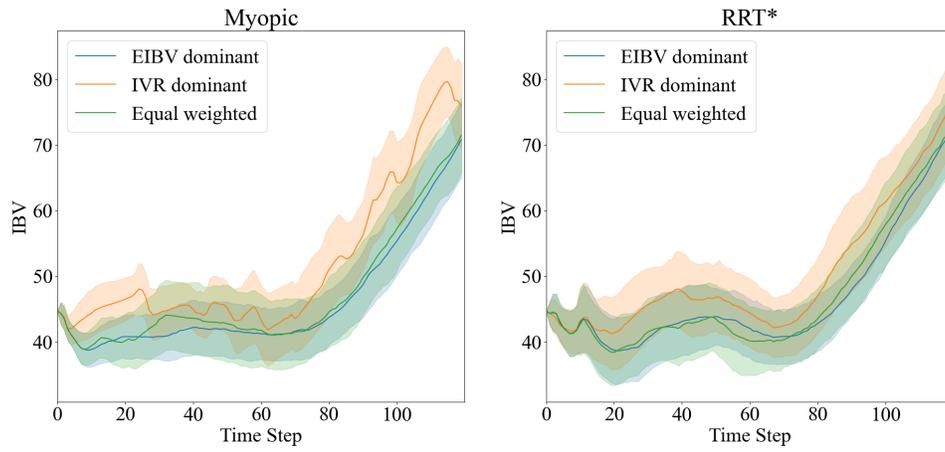


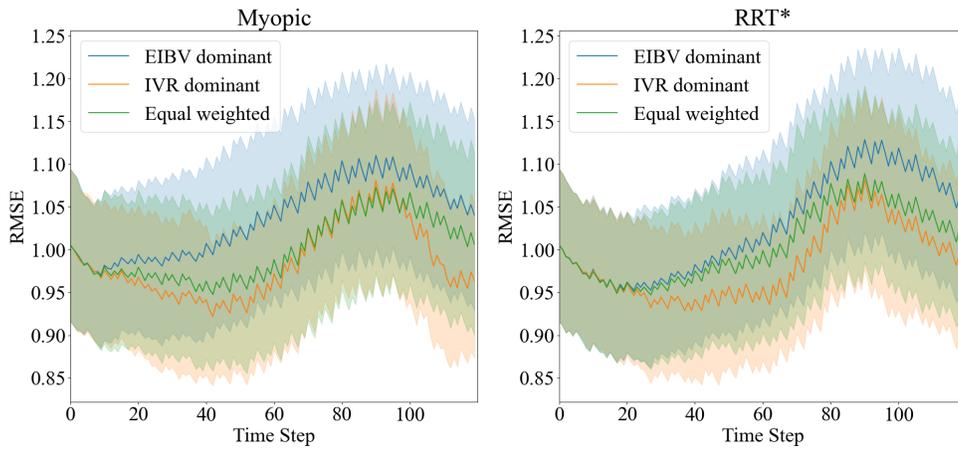
Figure 12: The traffic flow density map for the myopic agent (a) and the long-horizon RRT\* agent (b) using the equally weighted cost valley.

Figure 13 displays these metrics for the myopic and RRT\* planners across various weights in the cost valleys. Over time, these metrics vary. The IBV (top) increases for all variants during the later hours when the tide goes out and the plume front grows in extent. Naturally, the IVR dominant strategy has the poorest performance for this metric because it is not instructed to focus on EIBV. The equally weighted cost valley (green curves) exhibit balanced performance and is on par with the EIBV dominant strategy. RRT\* results of IBV tends to be slightly below the myopic strategy, but it is hardly significantly better for this metric. This is somewhat surprising after seeing the traffic analysis (Figure 8-12). It seems that RRT\* covers the domain better, but overall, when averaging over the entire grid, myopic performs similarly.

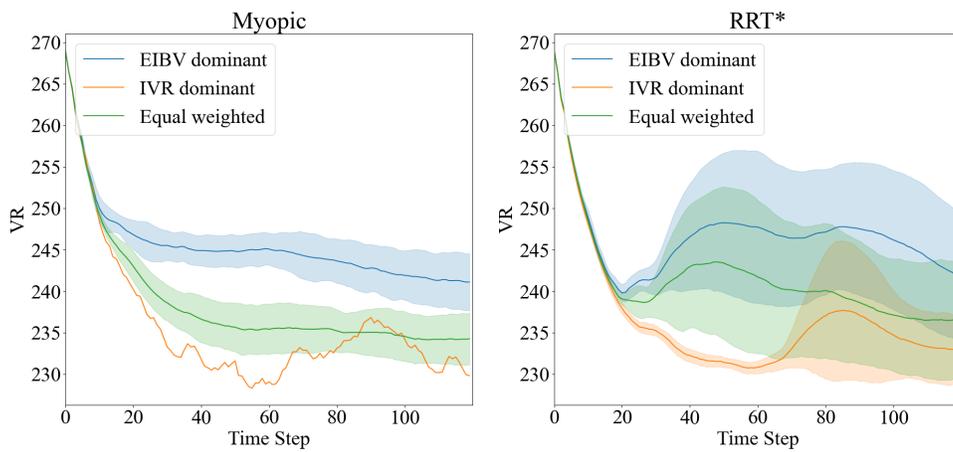
For the RMSE and VR metrics in Figure 13b-13c, the weighted strategy performs almost at the level of the IVR-dominant strategy which we would expect to do well at covering the domain and reducing variance and overall error compared with the truth, given the spatio-temporal correlation. Both weighted and IVR dominant appear substantially better than the EIBV dominant for these metrics. The differences between myopic and RRT\* strategies are again small. But for the VR metric, we see that all plots for RRT\* reach about 240 after 20 time steps, while the myopic strategies are not so good here. This is useful to know as it indicates that RRT\* with cost-valley can sample the field in a more efficient way. One does not always have 8 hours in the mission duration as that can reach the maximum battery.



(a)



(b)



(c)

Figure 13: The IBV (top), RMSE (middle) and VR (bottom) metrics from the 100 replicate simulation study. The solid line is the mean while the light-colored band indicates 90 % uncertainty over the 100 replicates.

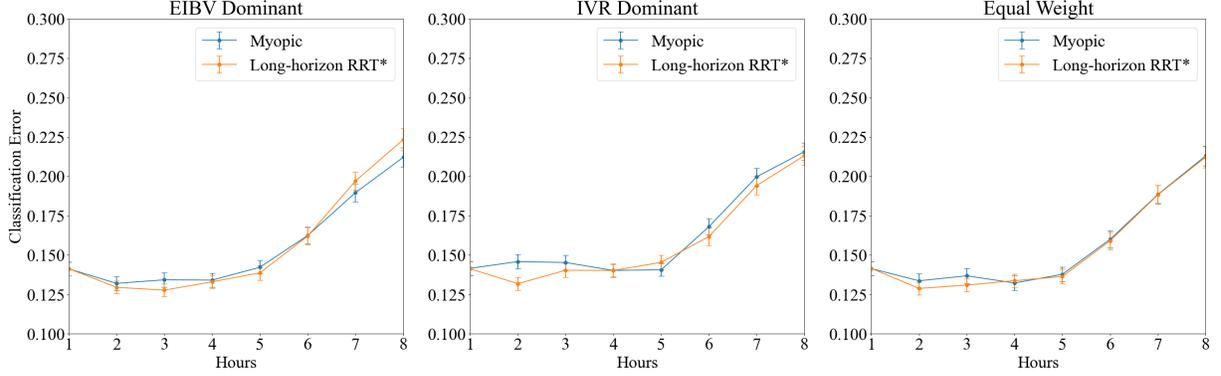


Figure 14: Classification error for the excursion sets over time using myopic or RRT\* strategies and three weighting schemes (left to right displays).

In Figure 14, we compare the classification error (CE) of each agent across different cost valleys. The CE for any salinity replicate realization is calculated using

$$\text{CE} = \frac{1}{B} \sum_{b=1}^B \frac{1}{n} \sum_{i=1}^n I[\text{ES}^{\text{true}}(t, \mathbf{u}_i) \neq \text{ES}^b(t, \mathbf{u}_i)], \quad (5)$$

where  $\text{ES}^b(t, \mathbf{u}_i)$  is the excursion set determined for each of  $b = 1, \dots, B$ ,  $B = 100$ , Monte Carlo samples from the updated Gaussian distribution at this time and  $\text{ES}^{\text{true}}(t, \mathbf{u}_i) = I(\xi_{t, \mathbf{u}_i}^{\text{true}} < \zeta)$  is the time- $t$  excursion set of the true replicate salinity field. This measure differs from the IBV in the sense that it goes beyond mere point-wise Bernoulli variance calculations, and it instead measures the uncertainty in the random sets. The CE in equation (5) is averaged over all replicates and evaluated every hour. Figure 14 shows that long-horizon RRT\* consistently gets a lower CE early in the operation, which is in line with our main goal. In a similar vein as for the IBV results, the CE grows over time because the river plume front goes out with the tide.

## 5.5 Remarks

Our simulation study shows the value of the cost valley in directing the agent. The EIBV cost field indicates probable river plume front locations, whereas the IVR cost field favors unexplored regions. By adjusting the weights of each cost field, we can refine the AUV's path planning and lose relatively little on the more focused metrics (Figure 13).

Traffic-flow analysis reveals the RRT\* planner's superiority over the myopic planner. The myopic planner's dependence on its starting point and its short-sightedness sometimes leads to prolonged confinement in

localized regions. In contrast, the RRT\* planner leverages the cost valley for broader and more efficient sampling, guiding the agent towards globally optimal paths.

Both myopic and RRT\* planner with cost valley manage to follow the tide over time, but the myopic strategy tends to lose one of the plume fronts because it gets too focused on one part for too long. The long-horizon RRT\* planner has been shown to have a slightly better performance in terms of its low classification error over time.

## 6 Experiment in the Trondheim fjord

We present the case of river plume exploration in the Trondheim fjord, Norway. The suggested algorithm using RRT\* with a cost valley runs onboard an AUV in this field experiment.

### 6.1 Experimental setup

The map view in Figure 2(a) displays the calibrated prior mean model for salinity within the spatial domain. The AUV runs its mission by utilizing Algorithm 2 for model updating and path planning based on the statistical surrogate model.

For this field experiment, we utilize a light autonomous underwater vehicle (LAUV) from NTNU’s applied underwater robotics laboratory (AURLab). Prior to launching the mission, the operator conducts a standard remote control check, as depicted in Figure 15.

The LAUV’s primary computing unit is the NVIDIA Jetson TX2. To enhance our onboard algorithm deployment capabilities, we use the adaptive sampling framework developed by Mo-Bjørkelund et al. (2020), which manages the messaging between ROS and DUNE. Our algorithm interfaces directly with Robot Operating Systems (ROS) (Quigley, 2009), and its messages are then relayed to the ROS-IMC bridge in the vehicle, which incorporates DUNE (DUNE: Unified Navigation Environment (Pinto et al., 2013)), as shown in Figure 16. The components within the LAUV communicate using the Inter Module Communication (IMC) message protocol (LSTS, 2022). We implement the same integration scheme as in Ge et al. (2023), which provides further details regarding the ROS-IMC bridge.

The LAUV is programmed to travel at a speed of 3 knots in the surface region at a depth of 0.5 m. Additionally, it is scheduled to resurface every 10 minutes to correct its navigational errors. The operation

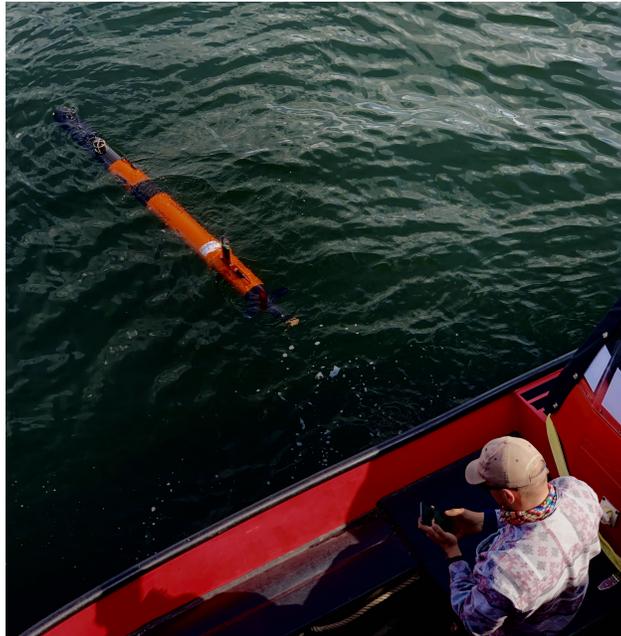


Figure 15: The LAUV named Thor is under remote operation check before its expedition.

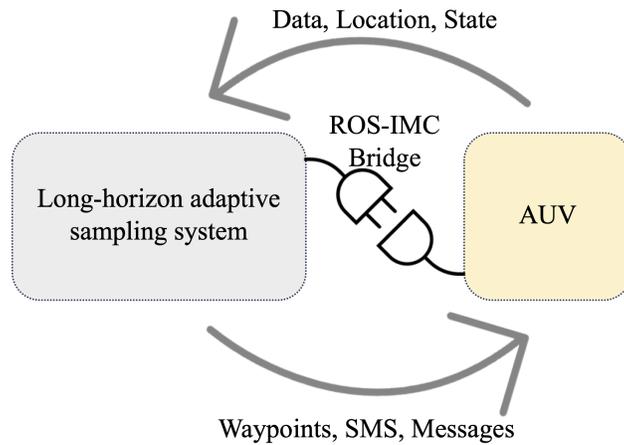


Figure 16: Main software components in the communication between the adaptive code and the vehicle. DUNE (Pinto et al., 2013) is running on the main CPU of the AUV while the IMC (LSTS, 2022) messages are transmitted via TCP (Cerf and Kahn, 1974) to an auxiliary CPU, where ROS (Quigley, 2009) and the adaptive code is run.

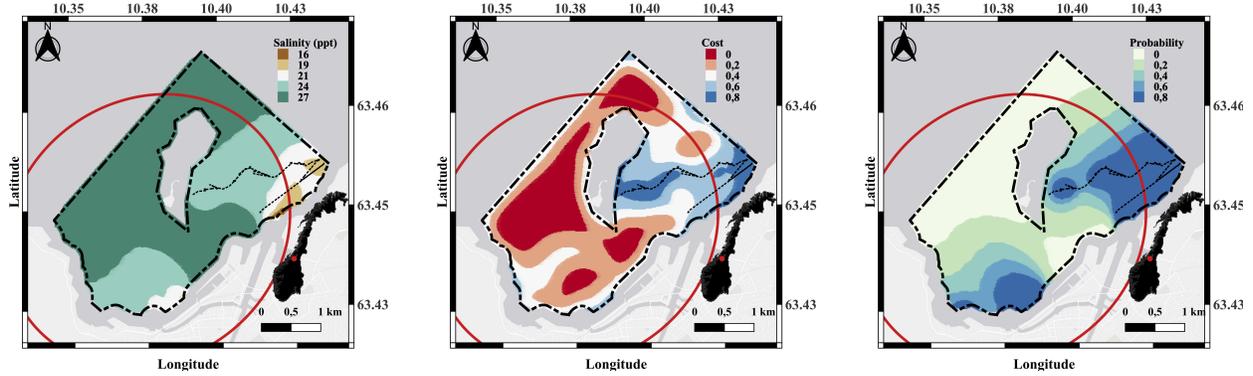


Figure 17: The updated mean field for the salinity (left), the cost valley (middle), and the excursion probability (right) after sampling the region from 12:30 to 13:41. The remaining budget starts to take an effect, shown as the red ellipse.

took place on May-11th, 2022, and it lasted for more than 2.5 hours. The LAUV left the start location at 12:30. We received the "Mission Complete" text message from the LAUV at 15:10, which marked the end of the operation.

## 6.2 Results

Figure 17 shows the updated mean salinity field (left display) after the LAUV has sampled data for 20-time steps (at 13:41). The associated cost valley field and excursion probability field are shown in Figure 17 (middle-right). The path taken by the LAUV (black line) indicates that it has used the first part of its adaptive mission to explore the area close to the river mouth.

As the AUV travels through the region, it adheres to the guidance provided by the cost valley and endeavors to minimize expenses. The budget ellipse (red color in Figure 17) diminishes as time goes by.

Figure 18 illustrates the same three spatial maps as in Figure 17, but now shown at step 30 (14:17). The AUV has skillfully avoided the obstacle in the middle. Furthermore, the map discloses the presence of a larger river plume in the western area, which contrasts the relatively small plume forecasted by SINMOD. At this time point, the agent's remaining budget is dwindling, as indicated by the red ellipse.

At step 44 (15:10), Figure 19 shows that the LAUV has discovered surprisingly low salinity values in the western parts. The excursion probability field (right display) also indicates a high probability of river water instead of ocean water quite far out to the north here. Currently, the available area enclosed by the red ellipse is restricted, and the LAUV must move along a relatively straight transect to reach the end location without much room for deviation.

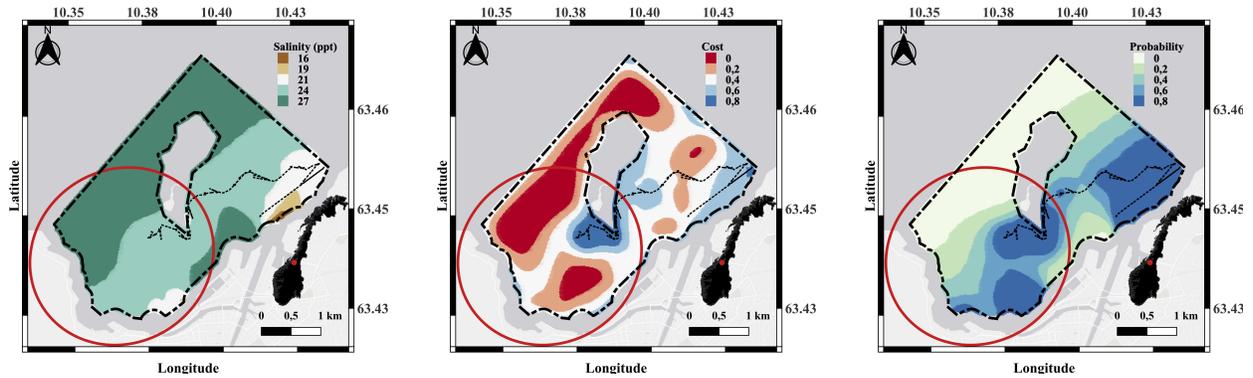


Figure 18: The updated salinity (left), cost (middle), and excursion probability (right) fields after collecting the in-situ measurements close to the obstacle area (step 30, at 14:17).

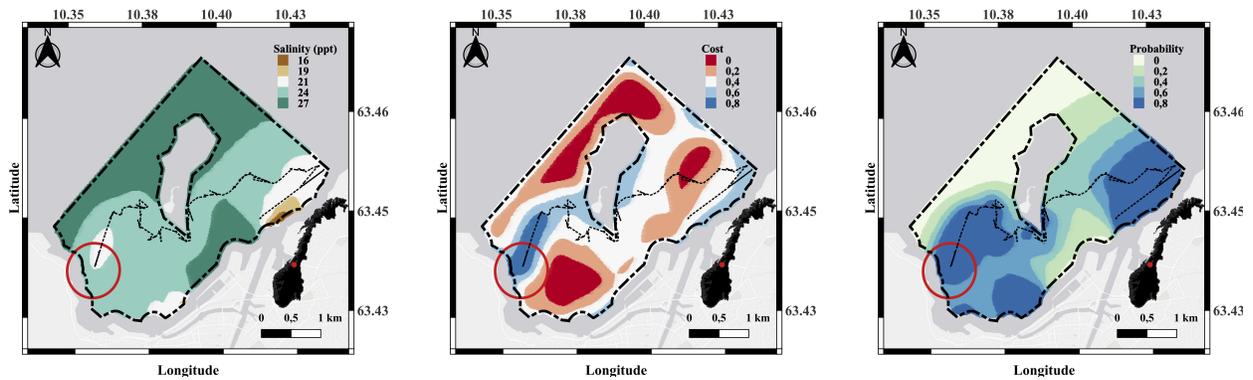


Figure 19: The revealed salinity field (left) together with the updated cost valley field (middle) and the excursion probability field (right) after that the LAUV sampled the north region using the last budget and safely returned home in the end (step 44, at 15:10).

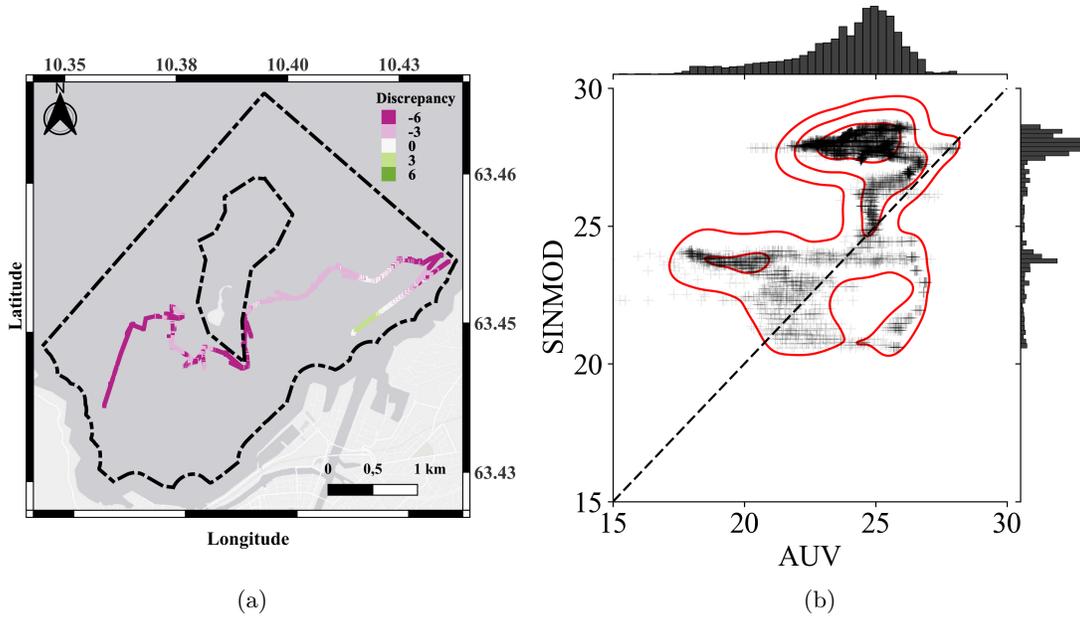


Figure 20: (a) The difference between the AUV measurements and the SINMOD prediction, referred to as the residual. The in-situ data reveal a considerable river plume in the western part, characterized by low salinity values, which contradicts the SINMOD prediction. (b) The cross plot between the SINMOD data and the AUV data. Kernel density estimation is applied here and it adds the red contours to the scatter plot. The histogram of each data source is shown on top and to the right.

Figure 20(a) shows the discrepancy between the in-situ AUV measurements and the SINMOD prediction. This is done by subtracting the SINMOD data at the AUV sampling locations from the in-situ measurements at these locations. These residuals are visualized in a map view and along the trajectory of the AUV. The plot has more negative than positive residuals, confirming that SINMOD tends to overestimate the salinity values in the region, resulting in a smaller river plume area. Figure 20(b) presents the cross plot of these two data sources. The kernel density estimate shows that SINMOD data has two modes near salinity 23 and 27, while the AUV data is distributed around salinity 25. The majority of the SINMOD data is higher than the actual measurements, indicating the need for practical adjustments of SINMOD data to avoid bias stemming from the numerical solver.

### 6.3 Remarks

The LAUV successfully navigated through the fjord waters, overcoming various constraints and acquiring sufficient information to reveal details of the river plume in the east as well as the one in the western region. This successful mission highlights the effectiveness of the suggested algorithm for ocean deployment. However, the use of auto-regressive modeling for the temporal aspect increases the system's dependence on

SINMOD, particularly in areas where samples were collected some time ago, resulting in greater uncertainty. While this approach is reasonable given the current mission duration, it may limit long-term performance.

## 7 Conclusion

The main contribution of this work is the development of a long-horizon adaptive sampling system using RRT\* path planning with a flexible cost valley in an informative field. We use this system in a field deployment with an AUV running autonomously from start location to home destination while adaptively sampling the salinity in a river plume front. In this field experiment, we used RRT\* as the path planner to determine the next waypoint and a budget ellipse for the time restriction consideration, all done in the context of the Gaussian random field model that is updated onboard the AUV. Using the cost valley concept for bridging multiple objective, the path planning achieves a balance between exploitation and exploration while the hard constraint on safety and punctuality are all considered and well shown in the final result of the field deployment.

In terms of the algorithm, an extensive simulation study shows the effect of weighting different objectives on the behavior of the AUV paths. In doing so, we further notice that the AUV changes its movement adaptively with the temporal dynamics of the river plume phenomenon. Comparing simulation results of a myopic strategy with that of long-horizon RRT\* path planning, we see that the myopic strategy is notably more focused on local details in the salinity map rather than the potential long-horizon benefits. Nevertheless, in terms of spatially integrated performance metrics such as root mean square error, the myopic strategy is not significantly worse than RRT\* in our situation.

In constructing the flexible cost valley, it is imperative that each constituent cost component exerts influence on the designated hot spot regions. As the number of objectives augments, there is an enhanced capability to address multiple areas of interest, contingent upon each component receiving an ample temporal allocation for utilization. A plethora of long-horizon path planning algorithms exist, inclusive of variants of RRT\*, A\* and the Probabilistic Road Map (PRM) could also be used here. The non-myopic capability of such planners facilitates the incorporation of future considerations within the sampling field, but it must run robustly in the contexts with computational or hardware constraints. We managed to run adaptive sampling with a long-horizon planner on a two-dimensional time-varying model. Adding depth variation would make it more computationally demanding. The system can be improved in the future by incorporating a dynamic obstacle avoidance system using the AIS system. Additionally, a more comprehensive temporal model can be added

to address the issue of relying too heavily on numerical data, as we currently use a standard auto-regressive model for temporal variation. Another potential improvement is the implementation of dynamic weighing mechanisms among the cost valley components to enhance its flexibility.

## Acknowledgments

We acknowledge support from the Norwegian Research Council (RCN) through the MASCOT project 305445. The authors thank AURLab NTNU for the support, collaboration, and easy access to testing equipment. We thank Tore Mo-Bjørkelund for his help during all the field trials. We thank SINTEF Ocean for supplying SINMOD data.

## References

- Anyosa, S., Eidsvik, J., and Pizarro, O. (2023). Adaptive spatial designs minimizing the integrated bernoulli variance in spatial logistic regression models-with an application to benthic habitat mapping. *Computational Statistics & Data Analysis*, 179:107643.
- Bai, S., Shan, T., Chen, F., Liu, L., and Englot, B. (2021). Information-driven path planning. *Current Robotics Reports*, 2(2):177–188.
- Belkin, I. M., Cornillon, P. C., and Sherman, K. (2009). Fronts in large marine ecosystems. *Progress in Oceanography*, 81(1-4):223–236.
- Berget, G. E., Fossum, T. O., Johansen, T. A., Eidsvik, J., and Rajan, K. (2018). Adaptive sampling of ocean processes using an auv with a gaussian proxy model. *IFAC-PapersOnLine*, 51(29):238–243. 11th IFAC Conference on Control Applications in Marine Systems, Robotics, and Vehicles CAMS 2018.
- Binney, J., Krause, A., and Sukhatme, G. S. (2013). Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8):873–888.
- Cerf, V. and Kahn, R. (1974). A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22(5):637–648.
- Cressie, N. and Wikle, C. K. (2015). *Statistics for spatio-temporal data*. John Wiley & Sons.
- Das, J., Rajany, K., Frolovy, S., Pyy, F., Ryany, J., Caronz, D. A., and Sukhatme, G. S. (2010). Towards

- marine bloom trajectory prediction for auv mission planning. In *2010 IEEE International Conference on Robotics and Automation*, pages 4784–4790. IEEE.
- Eidsvik, J., Mukerji, T., and Bhattacharjya, D. (2015). *Value of information in the earth sciences: Integrating spatial modeling and decision analysis*. Cambridge University Press.
- Enevoldsen, T. T. and Galeazzi, R. (2021). Grounding-aware rrt\* for path planning and safe navigation of marine crafts in confined waters. *IFAC-PapersOnLine*, 54(16):195–201.
- Ettlin, A. and Bleuler, H. (2006). Rough-terrain robot motion planning based on obstacleness. In *2006 9th International Conference on Control, Automation, Robotics and Vision*, pages 1–6.
- Fedorov, K. N. (1986). *The physical nature and structure of oceanic fronts*, volume 333. Springer.
- Fossum, T. O., Eidsvik, J., Ellingsen, I., Alver, M. O., Fragoso, G. M., Johnsen, G., Mendes, R., Ludvigsen, M., and Rajan, K. (2018). Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics*, 35(7):1101–1121.
- Fossum, T. O., Fragoso, G. M., Davies, E. J., Ullgren, J. E., Mendes, R., Johnsen, G., Ellingsen, I., Eidsvik, J., Ludvigsen, M., and Rajan, K. (2019). Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics*, 4(27):eaav3041.
- Fossum, T. O., Travelletti, C., Eidsvik, J., Ginsbourger, D., and Rajan, K. (2021). Learning excursion sets of vector-valued Gaussian random fields for autonomous ocean sampling. *The Annals of Applied Statistics*, 15(2):597 – 618.
- Ge, Y., Eidsvik, J., and Mo-Bjørkelund, T. (2023). 3d adaptive auv sampling for classification of water masses. *IEEE Journal of Ocean Engineering*, 48:626–639.
- Ge, Y., Olaisen, A. J. H., Eidsvik, J., Jain, R. P., and Johansen, T. A. (2022). Long-horizon informative path planning with obstacles and time constraints. *IFAC-PapersOnLine*, 55(31):124–129.
- Hollinger, G. A. and Sukhatme, G. S. (2014). Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287.
- Horner-Devine, A. R., Hetland, R. D., and MacDonald, D. G. (2015). Mixing and transport in coastal river plumes. *Annual Review of Fluid Mechanics*, 47:569–594.
- Hwang, J., Bose, N., and Fan, S. (2019). Auv adaptive sampling methods: A review. *Applied Sciences*, 9(15):3145.

- Jaillet, L., Cortés, J., and Siméon, T. (2010). Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics*, 26(4):635–646.
- Jakuba, M. and Yoerger, D. R. (2008). Autonomous search for hydrothermal vent fields with occupancy grid maps. In *Proc. of ACRA*, volume 8, page 2008. Citeseer.
- Karaman, S. and Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894.
- LSTS (2022). Inter module communication protocol.
- Ma, K.-C., Liu, L., Heidarsson, H. K., and Sukhatme, G. S. (2018). Data-driven learning and planning for environmental sampling. *Journal of Field Robotics*, 35(5):643–661.
- Mo-Bjørkelund, T., Fossum, T. O., Norgren, P., and Ludvigsen, M. (2020). Hexagonal grid graph as a basis for adaptive sampling of ocean gradients using auvs. In *Global Oceans 2020: Singapore – U.S. Gulf Coast*, pages 1–5.
- Okopal, G. (2019). Multi-objective autonomy for auv adaptive sampling missions. In *OCEANS 2019 MTS/IEEE SEATTLE*, pages 1–6. IEEE.
- Pinto, J., Dias, P. S., Martins, R., Fortuna, J., Marques, E., and Sousa, J. (2013). The lsts toolchain for networked vehicle systems. In *2013 MTS/IEEE OCEANS - Bergen*, pages 1–9.
- Preston, V., Flaspohler, G., Michel, A. P., Fisher III, J. W., and Roy, N. (2022). Robotic planning under uncertainty in spatiotemporal environments in expeditionary science. *arXiv preprint arXiv:2206.01364*.
- Quigley, M. (2009). Ros: an open-source robot operating system. In *ICRA 2009*.
- Richards, J. A. and Richards, J. A. (2022). *Remote sensing digital image analysis*, volume 5. Springer.
- Rigby, P., Pizarro, O., and Williams, S. B. (2010). Toward adaptive benthic habitat mapping using gaussian process classification. *Journal of Field Robotics*, 27(6):741–758.
- Sharples, J., Middelburg, J. J., Fennel, K., and Jickells, T. D. (2017). What proportion of riverine nutrients reaches the open ocean? *Global Biogeochemical Cycles*, 31(1):39–58.
- Slagstad, D. and McClimans, T. A. (2005). Modeling the ecosystem dynamics of the barents sea including the marginal ice zone: I. physical and chemical oceanography. *Journal of Marine Systems*, 58(1-2):1–18.
- Stankiewicz, P., Tan, Y. T., and Kobilarov, M. (2021). Adaptive sampling with an autonomous underwater vehicle in static marine environments. *Journal of Field Robotics*, 38(4):572–597.

- Suh, J., Gong, J., and Oh, S. (2017). Fast sampling-based cost-aware path planning with nonmyopic extensions using cross entropy. *IEEE Transactions on Robotics*, 33(6):1313–1326.
- Xiao, C. and Wachs, J. (2022). Nonmyopic informative path planning based on global kriging variance minimization. *IEEE Robotics and Automation Letters*, 7(2):1768–1775.
- Zacchini, L., Franchi, M., and Ridolfi, A. (2022). Sensor-driven autonomous underwater inspections: A receding-horizon rrt-based view planning solution for auvs. *Journal of Field Robotics*, 39(5):499–527.
- Zhang, Y., Bellingham, J. G., Ryan, J. P., Kieft, B., and Stanway, M. J. (2016). Autonomous four-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle. *Journal of Field Robotics*, 33(1):67–81.
- Zhang, Y., Kieft, B., Hobson, B. W., Ryan, J. P., Barone, B., Preston, C. M., Roman, B., Raanan, B.-Y., Marin III, R., O’Reilly, T. C., Rueda, C. A., Pargett, D., Yamahara, K. M., Poulos, S., Romano, A., Foreman, G., Ramm, H., Wilson, S. T., DeLong, E. F., Karl, D. M., Birch, J. M., Bellingham, J. G., and Scholin, C. A. (2020). Autonomous tracking and sampling of the deep chlorophyll maximum layer in an open-ocean eddy by a long-range autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 45(4):1308–1321.
- Zuhlke, M., Fomferra, N., Brockmann, C., Peters, M., Veci, L., Malik, J., and Regner, P. (2015). Snap (sentinel application platform) and the esa sentinel 3 toolbox. In *Sentinel-3 for Science Workshop*, volume 734, page 21.